

US008566928B2

(12) **United States Patent**
Dagon et al.

(10) **Patent No.:** **US 8,566,928 B2**
(45) **Date of Patent:** **Oct. 22, 2013**

(54) **METHOD AND SYSTEM FOR DETECTING AND RESPONDING TO ATTACKING NETWORKS**

(75) Inventors: **David Dagon**, Tampa, FL (US); **Nick Feamster**, Atlanta, GA (US); **Wenke Lee**, Atlanta, GA (US); **Robert Edmonds**, Woodstock, GA (US); **Richard Lipton**, Atlanta, GA (US); **Anirudh Ramachandran**, Atlanta, GA (US)

(73) Assignee: **Georgia Tech Research Corporation**, Atlanta, GA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1085 days.

(21) Appl. No.: **11/538,212**

(22) Filed: **Oct. 3, 2006**

(65) **Prior Publication Data**

US 2008/0028463 A1 Jan. 31, 2008

Related U.S. Application Data

(60) Provisional application No. 60/730,615, filed on Oct. 27, 2005, provisional application No. 60/799,248, filed on May 10, 2006.

(51) **Int. Cl.**
G06F 21/00 (2013.01)

(52) **U.S. Cl.**
USPC **726/22**

(58) **Field of Classification Search**
USPC **726/22, 23**
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,843,540	A	6/1989	Stolfo
4,860,201	A	8/1989	Stolfo et al.
5,363,473	A	11/1994	Stolfo et al.
5,497,486	A	3/1996	Stolfo et al.
5,563,783	A	10/1996	Stolfo et al.
5,668,897	A	9/1997	Stolfo
5,717,915	A	2/1998	Stolfo et al.
5,748,780	A	5/1998	Stolfo
5,920,848	A	7/1999	Schultzer et al.
6,401,118	B1	6/2002	Thomas
6,983,320	B1	1/2006	Thomas et al.
7,013,323	B1	3/2006	Thomas et al.
7,039,721	B1	5/2006	Wu et al.
7,069,249	B2	6/2006	Stolfo et al.

(Continued)

FOREIGN PATENT DOCUMENTS

WO	WO 02/37730	5/2002
WO	WO 02/098100	12/2002

OTHER PUBLICATIONS

“Spamming Botnets: Signatures and Characteristics” Xie et al; ACM SIGCOMM. Settle, WA; Aug. 2008; 12 pages.*

(Continued)

Primary Examiner — Taghi Arani

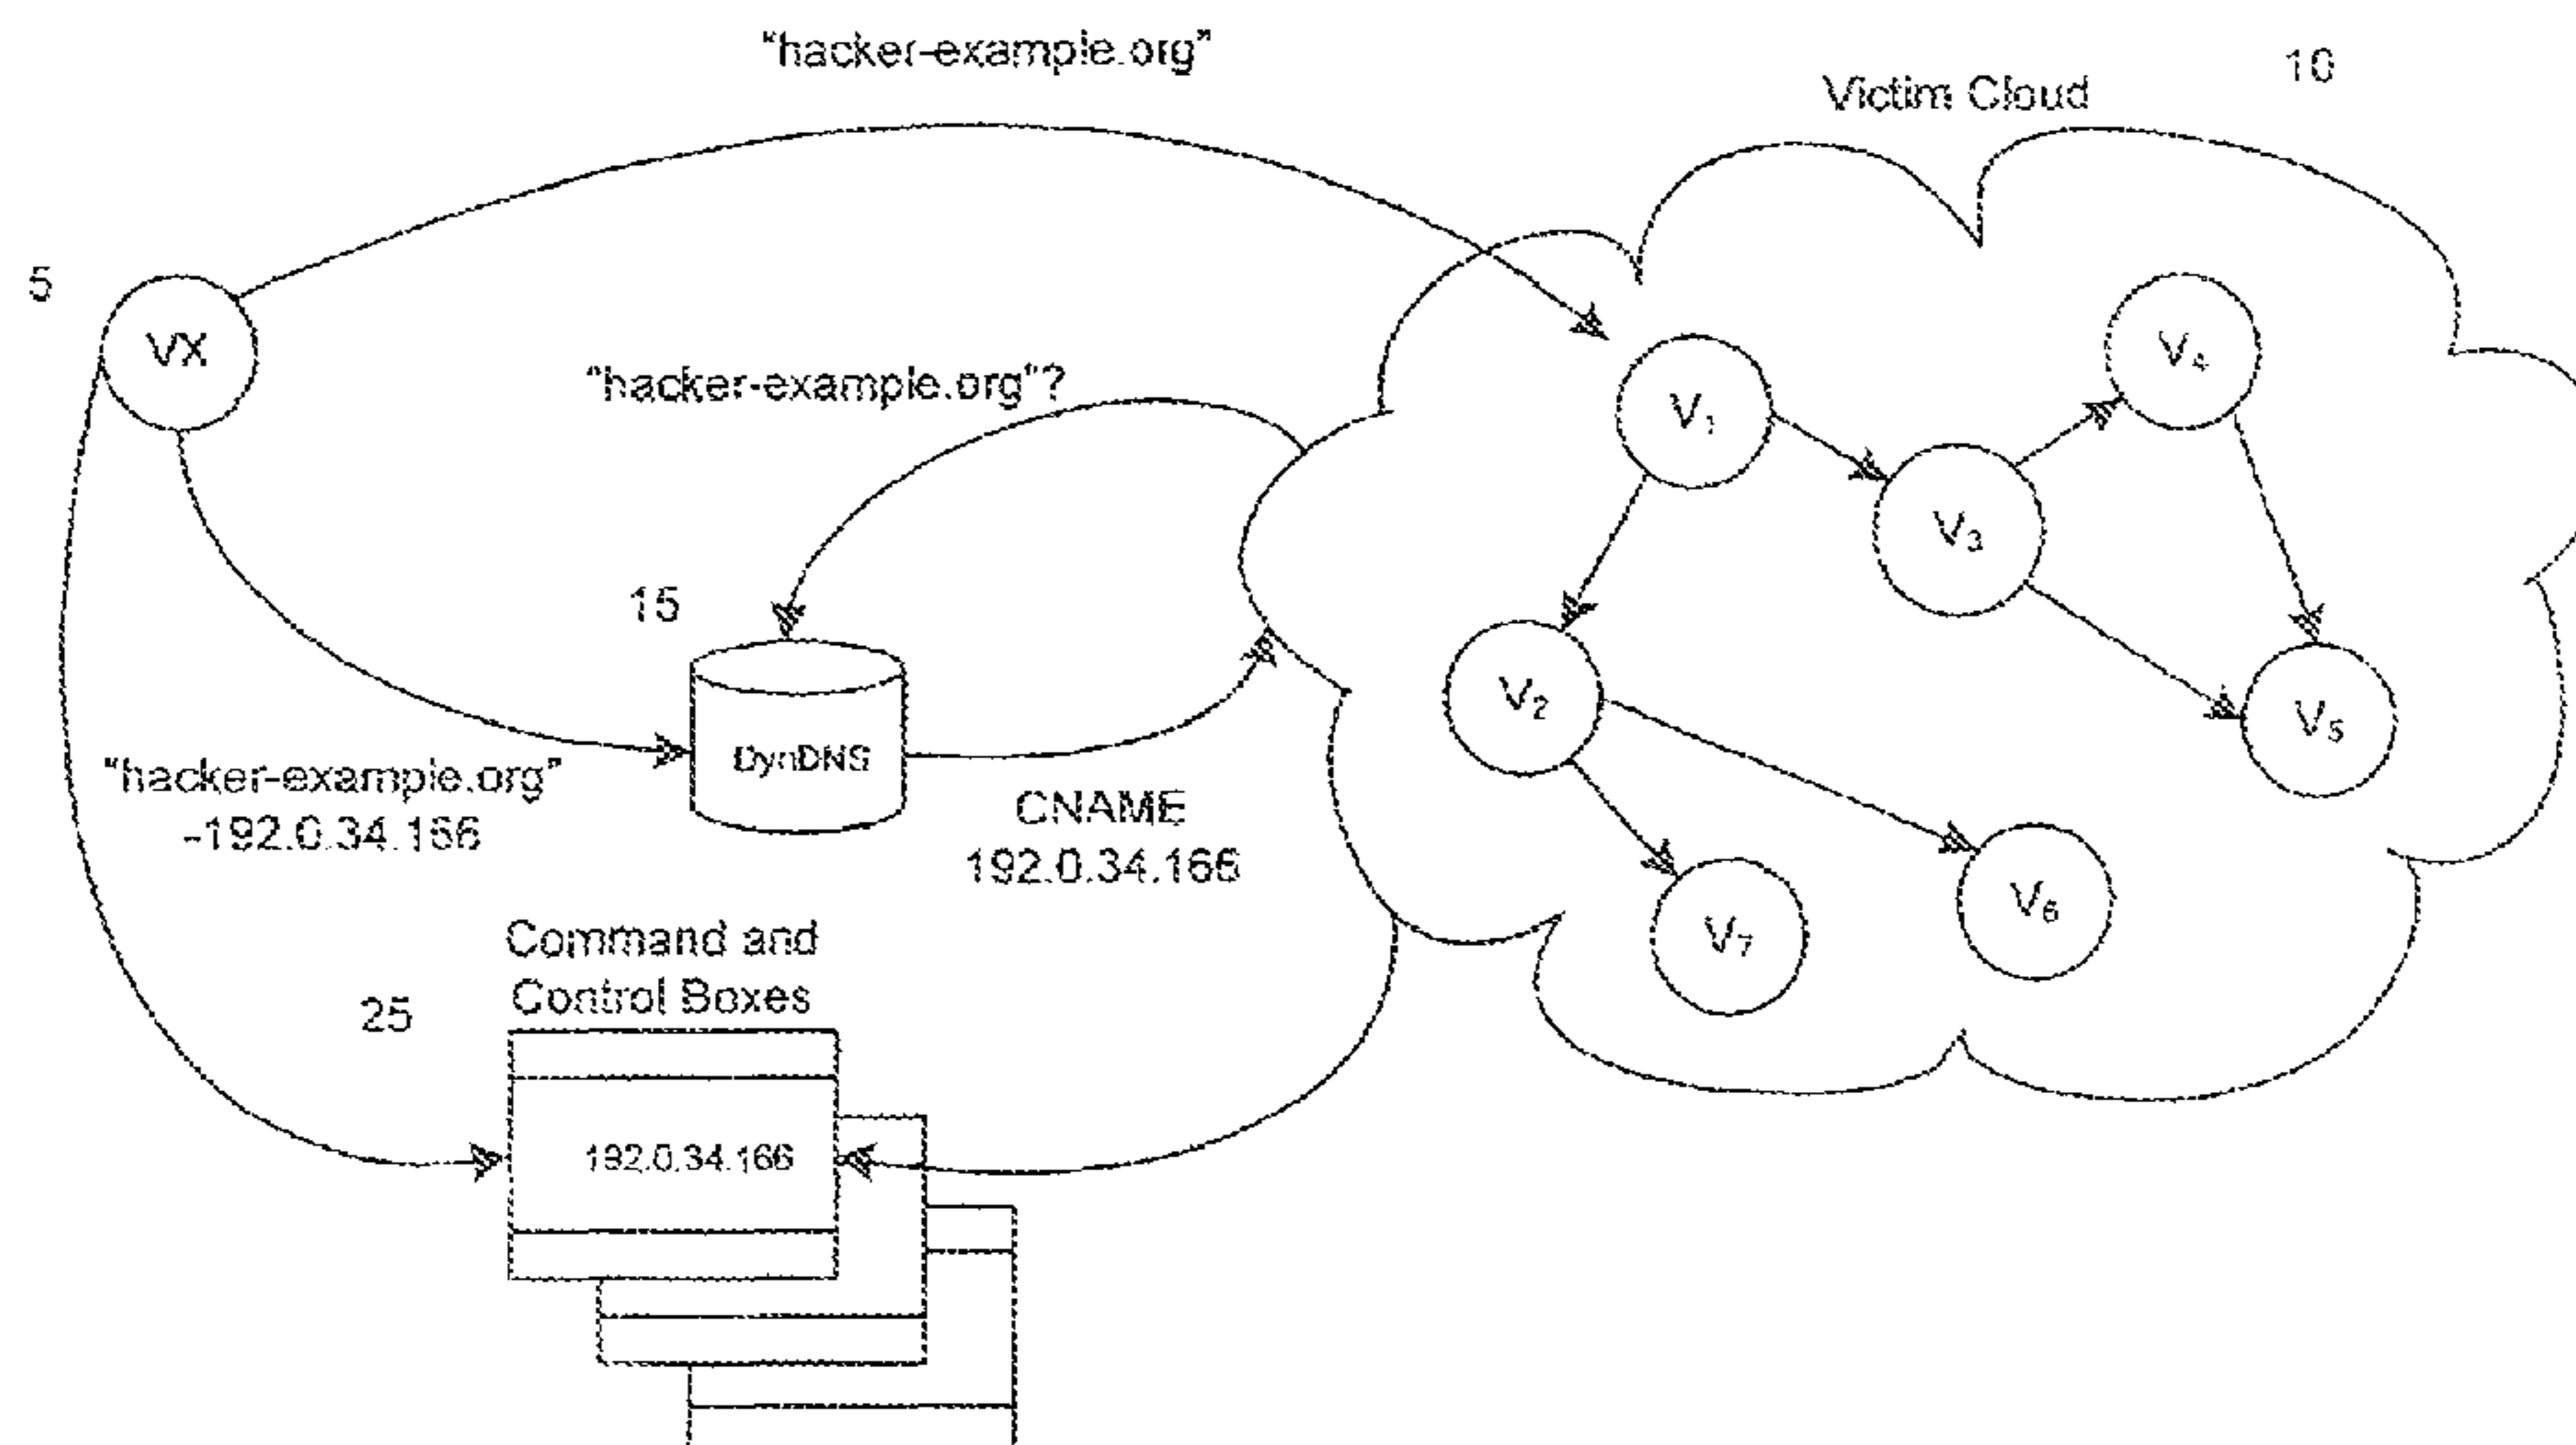
Assistant Examiner — Jason Lee

(74) *Attorney, Agent, or Firm* — DLA Piper LLP US

(57) **ABSTRACT**

A system and method for detecting a first network of compromised computers in a second network of computers, comprising: collecting Domain Name System (DNS) data for the second network; examining the collected data relative to DNS data from known comprised and/or uncompromised computers in the second network; and determining the existence of the first network and/or the identity of compromised computers in the second network based on the examination.

48 Claims, 28 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

7,093,292 B1 8/2006 Pantuso
 7,136,932 B1 11/2006 Schneider
 7,152,242 B2 12/2006 Douglas
 7,162,741 B2 1/2007 Eskin et al.
 7,225,343 B1 5/2007 Honig et al.
 7,277,961 B1 10/2007 Smith et al.
 7,331,060 B1 2/2008 Ricciulli
 7,372,809 B2 5/2008 Chen et al.
 7,383,577 B2* 6/2008 Hrastar et al. 726/23
 7,424,619 B1 9/2008 Fan et al.
 7,426,576 B1 9/2008 Banga et al.
 7,448,084 B1 11/2008 Apap et al.
 7,483,947 B2 1/2009 Starbuck
 7,487,544 B2 2/2009 Schultz et al.
 7,536,360 B2 5/2009 Stolfo et al.
 7,634,808 B1 12/2009 Szor
 7,639,714 B2 12/2009 Stolfo et al.
 7,657,935 B2 2/2010 Stolfo et al.
 7,665,131 B2 2/2010 Goodman
 7,698,442 B1 4/2010 Krishnamurthy
 7,752,125 B1 7/2010 Kothari et al.
 7,752,665 B1 7/2010 Robertson et al.
 7,779,463 B2 8/2010 Stolfo et al.
 7,784,097 B1 8/2010 Stolfo et al.
 7,818,797 B1 10/2010 Fan et al.
 7,913,306 B2 3/2011 Apap et al.
 7,930,353 B2 4/2011 Chickering
 7,962,798 B2 6/2011 Locasto et al.
 7,979,907 B2 7/2011 Schultz et al.
 7,996,288 B1 8/2011 Stolfo
 8,015,414 B2 9/2011 Mahone
 8,074,115 B2 12/2011 Stolfo et al.
 8,161,130 B2 4/2012 Stokes
 8,224,994 B1 7/2012 Schneider
 8,341,745 B1* 12/2012 Chau et al. 726/24
 2001/0044785 A1 11/2001 Stolfo et al.
 2001/0052007 A1 12/2001 Shigezumi
 2001/0052016 A1 12/2001 Skene et al.
 2001/0055299 A1 12/2001 Kelly
 2002/0021703 A1 2/2002 Tsuchiya et al.
 2002/0066034 A1 5/2002 Schlossberg et al.
 2003/0065926 A1 4/2003 Schultz et al.
 2003/0065943 A1 4/2003 Geis et al.
 2003/0069992 A1 4/2003 Ramig
 2003/0167402 A1 9/2003 Stolfo et al.
 2003/0204621 A1 10/2003 Poletto et al.
 2004/0002903 A1 1/2004 Stolfo et al.
 2004/0111636 A1 6/2004 Baffes et al.
 2004/0187032 A1 9/2004 Gels et al.
 2004/0205474 A1 10/2004 Eskin et al.
 2004/0215972 A1 10/2004 Sung et al.
 2005/0021848 A1* 1/2005 Jorgenson 709/238
 2005/0039019 A1 2/2005 Delany
 2005/0108407 A1 5/2005 Johnson et al.
 2005/0108415 A1 5/2005 Turk et al.
 2005/0257264 A1 11/2005 Stolfo et al.
 2005/0261943 A1 11/2005 Quarterman et al.
 2005/0265331 A1 12/2005 Stolfo
 2005/0281291 A1 12/2005 Stolfo et al.
 2006/0015630 A1 1/2006 Stolfo et al.
 2006/0075084 A1 4/2006 Lyon
 2006/0143711 A1* 6/2006 Huang et al. 726/23
 2006/0146816 A1* 7/2006 Jain 370/389
 2006/0156402 A1 7/2006 Stone et al.
 2006/0168024 A1 7/2006 Mehr
 2006/0178994 A1 8/2006 Stolfo et al.
 2006/0212925 A1 9/2006 Shull
 2006/0224677 A1 10/2006 Ishikawa et al.
 2006/0230039 A1 10/2006 Shull
 2006/0247982 A1 11/2006 Stolfo et al.
 2006/0253584 A1 11/2006 Dixon
 2007/0050708 A1 3/2007 Gupta et al.
 2007/0064617 A1 3/2007 Reves
 2007/0083931 A1 4/2007 Spiegel
 2007/0162587 A1 7/2007 Lund et al.

2007/0239999 A1 10/2007 Honig et al.
 2007/0274312 A1 11/2007 Salmela et al.
 2007/0294419 A1 12/2007 Ulevitch
 2008/0028073 A1 1/2008 Trabe et al.
 2008/0060054 A1 3/2008 Srivastava
 2008/0098476 A1 4/2008 Syversen
 2008/0155694 A1 6/2008 Kwon et al.
 2008/0229415 A1 9/2008 Kapoor
 2008/0276111 A1 11/2008 Jacoby et al.
 2009/0055929 A1 2/2009 Lee et al.
 2009/0083855 A1 3/2009 Apap et al.
 2009/0193293 A1 7/2009 Stolfo et al.
 2009/0222922 A1 9/2009 Sidiroglou et al.
 2009/0241191 A1 9/2009 Keromytis et al.
 2009/0254658 A1 10/2009 Kamikura et al.
 2009/0254992 A1 10/2009 Schultz et al.
 2010/0011243 A1 1/2010 Locasto et al.
 2010/0054278 A1 3/2010 Stolfo et al.
 2010/0064368 A1 3/2010 Stolfo et al.
 2010/0064369 A1 3/2010 Stolfo et al.
 2010/0138919 A1 6/2010 Peng
 2010/0146615 A1 6/2010 Locasto et al.
 2010/0169970 A1 7/2010 Stolfo et al.
 2010/0281541 A1 11/2010 Stolfo et al.
 2010/0281542 A1 11/2010 Stolfo et al.
 2011/0041179 A1 2/2011 Stahlberg
 2011/0214161 A1 9/2011 Stolfo et al.

OTHER PUBLICATIONS

Jelena Mirkovic et al., "Internet Denial of Service: Attack and Defense Mechanisms", pp. v-ix, 101-151, 153-220, 221-240 (2005).
 Joe Stewart, "Bobax Trojan Analysis", <http://www.lurhq.com/bobax.thml>, May 17, 2004.
 David Brumley et al., "Tracking Hackers on IRC", <http://www.doomed.com/texts/ircmirr/TrackingHackerson IRC.htm>, Dec. 8, 1999.
 Brian Krebs, "Bringing Botnet Out of the Shadows", http://www.washingtonpost.com/wp-dyn/content/article/2006/03/21/AR2006032100279_pf.html, Mar. 21, 2006.
 "SwatIT: Bots, Drones, Zombies, Worms and Other Things That Go Bump in the Night", <http://swatit.org/bots>, 2004.
 Christian Kreibich, "Honeycomb: Automated NIDS Signature Creation Using Honeypots", 2003, <http://www.cl.cam.ac.uk/research/srg/netos/papers/2003-honeycomb-sigcomm-poster.pdf>.
 DMOZ Open Directory Project, Dynamic DNS Providers List, http://dmoz.org/Computers/Software/Internet/Servers/Address_Management/Dynamic_DNS_Services/.
 David Moore, "Network Telescopes: Observing Small or Distant Security Events", http://www.caida.org/publications/presentations/2002/usenis_sec/usenix_sec_2002_files/frame.htm; Aug. 8, 2002.
 Vincent H. Berk et al., "Using Sensor Networks and Data Fusion for Early Detection of Active Worms", Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Defense and Law Enforcement II, Proceedings of SPIE, vol. 5071, pp. 92-104 (2003).
 International Search Report issued in Application No. PCT/US06/038611 mailed Jul. 8, 2008.
 Written Opinion issued in Application No. PCT/US06/038611 mailed Jul. 8, 2008.
 International Preliminary Report on Patentability issued in Application No. PCT/US06/038611 mailed Mar. 26, 2009.
 O. Diekmann et al., "Mathematical Epidemiology of Infectious Diseases: Model Building, Analysis and Interpretation", John Wiley & Son, Ltd., 2000, pp. v-xv and 1-303.
 Jelena Mirkovic et al., "Internet Denial of Service: Attack and Defense Mechanisms", Prentice Hall Professional Technical Reference, 2004, pp. v-xxii and 1-372.
 "Symantec Internet Security Threat Report: Trends for Jan. 1, 2004-Jun. 30, 2004" Symantec, Sep. 2004, pp. 1-54.
 David Dagon et al., "HoneyStat: Local Worm Detection Using Honeypots", RAID 2004, LNCS 3224, pp. 39-58 (2004).
 Jonghyun Kim et al., "Measurement and Analysis of Worm Propagation on Internet Network Topology", IEEE, pp. 495-500 (2004).

(56)

References Cited

OTHER PUBLICATIONS

- Andreas Marx, "Outbreak Response Times: Putting AV to the Test", www.virusbtn.com, Feb. 2004, pp. 4-6.
- Cliff Changchun Zou et al., "Worm Propagation Modeling and Analysis Under Dynamic Quarantine Defense", WORM'03, Oct. 27, 2003, Washington, DC USA, 10 pages.
- Thorsten Holz, "Anti-Honeypot Technology", 21st Chaos Communication Congress, slides 1-57, Dec. 2004.
- "CipherTrust's Zombie Stats", <http://www.ciphertrust.com/resources/statistics/zombie.php> 3 pages, printed Mar. 25, 2009.
- Joe Stewart, "Phatbot Trojan Analysis", <http://www.secureworks.com/research/threats/phatbot>, Mar. 15, 2004, 3 pages.
- Thorsten Holz et al., "A Short Visit to the Bot Zoo", IEEE Security & Privacy, pp. 76-79 (2005).
- Michael Glenn, "A Summary of DoS/DDoS Prevention, Monitoring and Mitigation Techniques in a Service Provider Environment", SANS Institute 2003, Aug. 21, 2003, pp. ii-iv, and 1-30.
- Dennis Fisher, "Thwarting the Zombies", Mar. 31, 2003, 2 pages.
- Felix C. Freiling et al., "Botnet Tracking: Exploring a Root-Cause Methodology to Prevent Distributed Denial-of-Service Attacks", ESORICS 2005, LNCS 3679, pp. 319-335 (2005).
- Vinod Yegneswaran et al., "Global Intrusion Detection in the DOMINO Overlay System", Proceedings of Network and Distributed Security Symposium (NDSS), 17 pages Feb. 2004.
- Vinod Yegneswaran et al., "On the Design and Use of Internet Sinks for Network Abuse Monitoring", RAID 2004, LNCS 3224, pp. 146-165 (2004).
- Cliff C. Zou et al., "Routing Worm: A Fast, Selective Attack Worm Based on IP Address Information", Technical Report: TR-03-CSE-06, Principles of Advanced and Distributed Simulation (PADS) 2005, pp. 199-206, Jun. 1-3, 2005.
- File History for U.S. Appl. No. 12/538,612, filed Aug. 10, 2009 (downloaded Apr. 29, 2010).
- Dongeun Kim et al., "Request Rate Adaptive Dispatching Architecture for Scalable Internet Server", Proceedings of the IEEE International Conference on Cluster Computing (CLUSTER'00); pp. 289-296 (2000).
- Keisuke Ishibashi et al., "Detecting Mass-Mailing Worm Infected Hosts by Mining DNS Traffic Data", SIGCOMM'05 Workshops, pp. 159-164 (Aug. 22-26, 2005).
- Nicholas Weaver et al., "A Taxonomy of Computer Worms", WORM'03, pp. 11-18 (Oct. 27, 2003).
- File History of U.S. Appl. No. 11/538,212.
- Stephan Axelsson, "The Base-Rate Fallacy and the Difficulty of Intrusion Detection", ACM Transactions on Information and System Security, vol. 3, No. 3, pp. 186-205 (Aug. 2000).
- Niel Landwehr et al., "Logistic Model Trees", Machine Learning, vol. 59, pp. 161-205 (2005).
- Richard O. Duda et al., "Pattern Classification, Second Edition", John Wiley & Sons, Inc., pp. vii-xx, and 1-654, Copyright 2001.
- File History of U.S. Appl. No. 12/538,612.
- File History of U.S. Appl. No. 12/985,140.
- File History of U.S. Appl. No. 13/008,257.
- File History of U.S. Appl. No. 13/205,928.
- File History of U.S. Appl. No. 13/309,202.
- File History of U.S. Appl. No. 13/358,303.
- File History of U.S. Appl. No. 13/749,205.
- P. Mockapetris, "Domain Names—Concepts and Facilities", Network Working Group, <http://www.ietf.org/rfc/rfc1034.txt>, Nov. 1987 (52 pages).
- P. Mockapetris, "Domain Names—Implementation and Specification", Network Working Group, <http://www.ietf.org/rfc/rfc1035.txt>, Nov. 1987 (52 pages).
- P. Akritidis et al., "Efficient Content-Based Detection of Zero-Day Worms", 2005 IEEE International Conference in communications, vol. 2, pp. 837-843, May 2005.
- Nicholas Weaver et al., "Very Fast Containment of Scanning Worms", In proceedings of the 13th USENIX Security in Symposium, pp. 29-44, Aug. 9-13, 2004.
- David Whyte et al., "DNS-Based Detection of Scanning Worms in an Enterprise Network", In Proc. of the 12th Annual Network and Distributed System Security Symposium, pp. 181-195, Feb. 3-4, 2005.
- Cristian Abad et al., "Log Correlation for Intrusion Detection: A Proof of Concept", In Proceedings of The 19th Annual Computer Security Application Conference (ACSAC'03), (11 pages) (2003).
- Lala A. Adamic et al., "Zipf's Law and the Internet", Glottometrics, vol. 3, pp. 143-150 (2002).
- K.G. Anagnostakis et al., "Detecting Targeted Attacks Using Shadow Honeypots", In Proceedings of the 14th USENIX Security Symposium, pp. 129-144 (2005).
- Paul Baecher et al., "The Nepenthes Platform: An Efficient Approach to Collect Malware", In Proceedings of Recent Advances in Intrusion Detection (RAID 2006), LNCS 4219, pp. 165-184, Sep. 2006.
- Paul Barford et al., "An Inside Look at Botnets", Special Workshop on Malware Detection, Advances in Information Security, Springer Verlag, pp. 171-192 (2006).
- James R. Binkley et al., "An Algorithm for Anomaly-Based Botnet Detection", 2nd Workshop on Steps to Reducing Unwanted Traffic on the Internet (SRUTI '06), pp. 43-48, Jul. 7, 2006.
- Steven Cheung et al., "Modeling Multistep Cyber Attacks for Scenario Recognition", In Proceedings of the Third DARPA Information Survivability Conference and Exposition (DISCEX III), vol. 1, pp. 284-292, Apr. 22-24, 2003.
- Evan Cooke et al., "The Zombie Roundup: Understanding, Detecting, and Disrupting Botnets", Steps to Reducing Unwanted Traffic on the Internet Workshop (SRUTI '05), pp. 39-44, Jul. 7, 2005.
- Frederic Cuppens et al., "Alert Correlation in a Cooperative Intrusion Detection Framework", In Proceedings of IEEE Symposium on Security and Privacy 2002, pp. 202-215 (2002).
- David Dagon et al., "Modeling Botnet Propagation using Time Zones", The 13th Annual Network and Distributed System Security Symposium 2006, Feb. 2-3, 2006 (18 pages).
- Roger Dingledine et al., "Tor: The Second-Generation Onion Router", In Proceedings of the 13th Usenix Security Symposium, pp. 303-320 Aug. 9-13, 2004.
- Steven T. Eckman et al., "STATL: An Attack Language for State-Based Intrusion Detection", Journal of Computer Security, vol. 10, pp. 71-103 (2002).
- Daniel R. Ellis, et al., "A Behavioral Approach to Worm Detection", WORM'04, Oct. 29, 2004 (11 pages).
- Prahlad Fogla et al., "Polymorphic Blending Attacks", In Proceedings of 15th Usenix Security Symposium, pp. 241-256, (2006).
- Koral Ilgun et al., "State transition Analysis: A Rule-Based Intrusion Detection Approach", IEEE Transactions on Software Engineering, vol. 21, No. 3, pp. 181-199, Mar. 1995.
- Giovanni Vigna et al., "NetSTAT: A Network-based Intrusion Detection Approach", In Proceedings of the 14th Annual Computer Security Applications Conference (ACSAC '98), pp. 25-34, Dec. 7-11, 1998.
- Christopher Kruegel et al., "Polymorphic Worm Detection using Structural Information of Executables", RAID 2005, pp. 207-226 (2005).
- Ke Wang et al., "Anagram: A Content Anomaly Detector Resistant to Mimicry Attack", In Proceedings of the International Symposium on Recent Advances in Intrusion Detection (RAID) (2006) (20 pages).
- Ke Wang et al., "Anomalous Payload-Based Worm Detection and Signature Generation", In Proceedings of the International Symposium on Recent Advances in Intrusion Detection (RAID) (2005) (20 pages).
- David Whyte, "Exposure Maps: Removing Reliance on Attribution During Scan Detection", 1st Usenix Workshop on Hot Topics in Security, pp. 51-55 (2006).
- Jiahai Yang et al., "CARDS: A Distributed System for Detecting Coordinated Attacks", In Sec (2000) (10 pages).
- Vinod Yegneswaran et al., "Using Honeynets for Internet Situational Awareness", In proceedings of the Fourth Workshop on Hot Topics in Networks (HotNets IV), Nov. 2005 (6 pages).
- Jaeyeon Jung et al., "DNS Performance and the Effectiveness of Caching", IEEE/ACM Transactions on Networking, vol. 10, No. 5, pp. 589-603, Oct. 2002.
- Duane Wessels et al., "Measurements and Laboratory Simulations of the Upper DNS Hierarchy", In PAM (2005) (10 pages).

(56)

References Cited

OTHER PUBLICATIONS

- Paul Barham et al., "Xen and the Art of Virtualization", SOSP'03, Oct. 19-22, 2003 (14 pages).
- Ulrich Bayer et al., "TTAnalyze: A Tool for Analyzing Malware", In Proceedings of the 15th Annual Conference European Institute for Computer Antivirus Research(EICAR), pp. 180-192 (2006).
- Fabrice Bellard, "QEMU, A Fast and Portable Dynamic Translator", In Proceedings of the Annual Conference on Usenix Annual Technical Conference, pp. 41-46 (2005).
- Kevin Borders et al., "Siren: Catching Evasive Malware (Short Paper)", IEEE Symposium on Security and Privacy, pp. 78-85, May 21-24, 2006.
- Christopher M. Bishop, Pattern Recognition and Machine Learning (Information Science and Statistics), Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.
- Matt Bishop, "Computer Security: Art and Science", Addison-Wesley Professional, 2003.
- Michael Sipser, "Introduction to the Theory of Computation", International Thomson Publishing, 1996.
- Peter Szor, "The Art of Computer Virus Research and Defense", Addison-Wesley Professional, 2005.
- Anil K. Jain et al., "Algorithms for Clustering Data", Prentice-Hall, Inc., 1988.
- V. Laurikari, "TRE", 2006 (5 pages).
- Changda Wang et al., "The Dilemma of Covert Channels Searching", ICISC 2005, LNCS 3935, pp. 169-174, 2006.
- Mihai Christodorescu et al., "Semantics-Aware Malware Detection", In Proceeding of the 2005 IEEE Symposium on Security and Privacy, pp. 32-46 (2005).
- Peter Ferrie, "Attacks on Virtual Machine Emulators", Symantec Advance Threat Research, 2006 (13 pages).
- Tal Garfinkel et al., "A Virtual Machine Introspection Based Architecture for Intrusion Detection", In Proceedings of Network and Distributed Systems Security Symposium, Feb. 2003 (16 pages).
- G. Hunt et al., "Detours: Binary Interception of WIN32 Functions", Proceedings of the 3rd Usenix Windows NT Symposium, Jul. 12-13, 1999 (9 pages).
- Xuxian Jiang et al., "Virtual Playgrounds for Worm Behavior Investigation", RAID 2005, LNCS 3858, pp. 1-21 (2006).
- Christopher Kruegel et al., "Detecting Kernel-Level Rootkits Through Binary Analysis", In Proceedings of the Annual Computer Security Applications Conference (ACSAC), pp. 91-100, Dec. 2004.
- Paul Royal et al., "PolyUnpack: Automating the Hidden-Code Extraction of Unpack-Executing Malware", In Proceedings of the Annual Computer Security Applications Conference (ACSAC), pp. 289-300 (2006).
- Rich Uhlig et al., "Intel Virtualization Technology", Computer, vol. 38, No. 5, pp. 48-56, May 2005.
- Amit Vasudevan et al., "Stealth Breakpoints", In Proceedings of the 21st Annual Computer Security Applications Conference (ACSAC), pp. 381-392, (2005).
- Amit Vasudevan et al., "Cobra: Fine-Grained Malware Analysis Using Stealth Localized-Executions", In Proceedings of the 2006 IEEE Symposium on Security and Privacy (S&P'06), pp. 264-279 (2006).
- Yi-Min Wang et al., "Automated Web Patrol with Strider HoneyMonkeys: Finding Web Sites That Exploit Browser Vulnerabilities", In NDSS'06 (2006) (15 pages).
- Joanna Rutkowska, "Introducing Blue Pill", <http://theinvisiblethings.blogspot.com/2006/06/introducing-blue-pill.html>, Jun. 22, 2006 (26 pages).
- Maria Halkidi et al., "On Clustering Validation Techniques", Journal of Intelligent Information Systems, vol. 17, pp. 107-145 (2001).
- A.K. Jain et al., "Data Clustering: A Review", ACM Computing Surveys, vol. 31, No. 3, pp. 264-323, Sep. 1999.
- Hyang-Ah Kim et al., "Autograph: Toward Automated, distributed Worm Signature Detection", In Usenix Security Symposium (2004) (16 pages).
- Christian Kreibich et al., "Honeycomb—Creating Intrusion Detection Signatures Using Honey Pots", In ACM Workshop on Hot Topics in Networks (2003) (6 pages).
- Zhichun Li et al., "Hamsa: Fast Signature Generation for Zero-Day Polymorphic Worms with Provable Attack Resilience", In IEEE Symposium on Security and Privacy (2006) (15 pages).
- James Newsome et al., "Polygraph: Automatically Generating Signatures for Polymorphic Worms", In IEEE Symposium on Security and Privacy (2005) (16 pages).
- Sun Wu et al., "AGREP—A Fast Approximate Pattern-Matching Tool", In Usenix Technical Conference (1992) (10 pages).
- Vinod Yegneswaran et al., "An Architecture for Generating Semantics-Aware Signatures", In Usenix Security Symposium (2005) (16 pages).
- Jaeyeon Jung, "Fast Portscan Detection Using Sequential Hypothesis Testing", In Proceedings of IEEE Symposium on Security Privacy, pp. 211-225 (2004).
- Carl Livades et al., "Using Machine Learning Techniques to Identify Botnet Traffic", In 2nd IEEE LCN Workshop on Network Security (WoNS'2006), pp. 967-974 (2006).
- "CVE-2006-3439", <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2006-3439>, printed Jun. 27, 2012 (2 pages).
- David Moore, "Inferring Internet Denial-of-Service Activity", In Proceedings of the 10th Usenix Security Symposium, Aug. 13-17, 2001 (15 pages).
- Peng Ning et al., "Constructing Attack Scenarios Through Correlation of Intrusion Alerts", In Proceedings of Computer and Communications Security (CCS'02), Nov. 18-22, 2002 (10 pages).
- Vern Paxson, "Bro: A System for Detecting Network Intruders in Real-Time", In Proceedings of the 7th Usenix Security Symposium, Jan. 26-29, 1998 (22 pages).
- Phillip A. Porras, "Privacy-Enabled Global Threat Monitoring", IEEE Security & Privacy, pp. 60-63 (2006).
- Anirudh Ramachandran et al., "Understanding the Network-Level Behavior of Spammers", In Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM'06), Sep. 11-16, 2006 (13 pages).
- Martin Roesch, "SNORT—Lightweight Intrusion Detection for Networks", In Proceedings of 13th System Administration Conference (LISA'99), pp. 229-238, Nov. 7-12, 1999.
- Robin Sommer et al., "Enhancing Byte-Level Network Intrusion Detection Signatures with Context", In Proceedings of the 10th ACM Conference on Computer and Communications Security (CCS'03), pp. 262-271, Oct. 27-30, 2003.
- Stuart Staniford et al., "Practical Automated Detection of Stealthy Portscans", Journal of Computer Security, vol. 10, pp. 105-136 (2002).
- S. Staniford-Chen et al., "GrIDS—A Graph Based Intrusion Detection System for Large Networks", In Proceedings of the 19th National Information Systems Security Conference, pp. 361-370 (1996).
- Steven J. Templeton et al., "A Requires/Provides Model for Computer Attacks", In Proceedings of the 2000 Workshop on New Security Paradigms (NSPW'00), pp. 31-38 (2000).
- Alfonso Valdes et al., "Probabilistic Alert Correlation", In Proceedings of the Recent Attack in Intrusion Detection (RAID 2001), LNCS 2212, pp. 54-68 (2001).
- Fredrik Valeur et al., "A Comprehensive Approach to Intrusion Detection Alert Correlation", IEEE Transactions on Dependable and Secure Computing, vol. 1, No. 3, pp. 146-169, Jul. 2004.
- Kjersti Aas et al., "Text Categorisation: A Survey", Norwegian Computing Center, Jun. 1999 (38 pages).
- M. Andrews, "Negative Caching of DNS Queries (DNS NCACHE)", <http://tools.ietf.org/html/rfc2308>, Mar. 1998 (20 pages).
- Simon Biles, "Detecting the Unknown with Snort and Statistical Packet Anomaly Detecting Engine", www.cs.luc.edu/~pld/courses/447/sum08/class6/biles.spade.pdf (2003) (9 pages).
- James Newsome et al., "Paragraph: Thwarting Signature Learning by Training Maliciously", In Recent Advance in Intrusion Detection (RAID), 2005 (21 pages).

(56)

References Cited

OTHER PUBLICATIONS

Dan Pelleg et al., "X-Means: Extending K-Means with Efficient Estimation of the Number of Clusters", In International Conference on Machine Learning (2000) (8 pages).

Roberto Perdisci et al., "Misleading Worm Signature Generators Using Deliberate Noise Injection", In IEEE Symposium on Security and Privacy (2006) (15 pages).

Sumeet Singh et al., "Automated Worm Fingerprinting", In ACM/USENIX Symposium on Operating System Design and Implementation, Dec. 2004 (16 pages).

R. Arends et al., "Protocol Modifications for the DNS Security Extensions", <http://www.ietf.org/rfc/rfc4035.txt>, Mar. 2005 (50 pages).

R. Arends et al., "DNS Security Introduction and Requirements", <http://www.ietf.org/rfc/rfc4033.txt>, Mar. 2005 (20 pages).

R. Arends et al., "Resource Records for the DNS Security Extensions", <http://www.ietf.org/rfc/rfc4034.txt>, Mar. 2005 (28 pages).

Jaeyeon Jung et al., "Modeling TTL-Based Internet Caches", IEEE INFOCOM 2003, pp. 417-426, Mar. 2003.

Florian Weimer, "Passive DNS Replication", In Proceedings of the 17th Annual FIRST Conference on Computer Security Incident, Apr. 2005 (13 pages).

Manos Antonakakis et al., "Unveiling the Network Criminal Infrastructure of TDSS/TDL4", http://www.damballa.com/downloads/r_pubs/Damballa_tdss_tdl4_case_study_public.pdf, (undated) (16 pages).

T. Berners-Lee et al., "RFC3986—Uniform Resource Identifier (URI): Generic Syntax", <http://www.hjp.at/doc/rfc/rfc3986.html>, Jan. 2005 (62 pages).

D. De La Higuera et al., "Topology of Strings: Median String is NP-Complete", Theoretical Computer Science, vol. 230, pp. 39-48 (2000).

John C. Platt, "Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods", Advances in Large margin Classifiers, vol. 10, No. 3, pp. 61-74, Mar. 26, 1999.

Nello Cristianini et al., "An Introduction to Support Vector Machines: and other Kernel-Based Learning Methods", Cambridge University Press, New York, NY, USA (2000).

<http://www.bleedingsnort.com>, retrieved from Internet Archive on May 23, 2013, Archived Sep. 26, 2006 (3 pages).

<http://www.dshield.org>, retrieved from Internet Archive on May 23, 2013, Archived Sep. 29, 2006 (2 pages).

<http://www.alexandria.com>, retrieved from Internet Archive on May 23, 2013, Archived Sep. 25, 2006 (3 pages).

<http://www.dnswl.org>, retrieved from Internet Archive on May 23, 2013, Archived Jul. 15, 2006 (4 pages).

<http://www.spamhaus.org/sbl/>, retrieved from Internet Archive on May 23, 2013, Archived Sep. 24, 2006 (24 pages).

<http://www.opendns.com>, retrieved from Internet Archive on May 23, 2013, Archived Sep. 9, 2006 (25 pages).

<http://www.avira.com>, retrieved from Internet Archive on May 23, 2013, Archived Sep. 29, 2006 (13 pages).

<http://www.oreans.com/themida.php>, retrieved from Internet Archive on May 23, 2013, Archived Aug. 23, 2006 (12 pages).

<http://www.vmware.com>, retrieved from Internet Archive on May 23, 2013, Archived Sep. 26, 2006 (32 pages).

<http://www.siliconrealms.com>, retrieved from Internet Archive on May 23, 2013, Archived Sep. 4, 2006 (12 pages).

<http://www.dyninst.org>, retrieved from Internet Archive on May 23, 2013, Archived Aug. 20, 2006 (pages).

F. Heinz et al., "IP Tunneling Through Nameserver", <http://slashdot.org/story/00/09/10/2230242/ip-tunneling-through-nameservers>, Sep. 10, 2000 (23 Pages).

Zhuoqing Morley Mao et al., "A Precise and Efficient Evaluation of the Proximity Between Web Clients and Their Local DNS Servers", In Proceedings of USENIX Annual Technical Conference (2002) (14 pages).

Cliff Changchun Zou et al., "Code Red Worm Propagation Modeling and Analysis", In Proceedings of 9th ACM Conference on Computer and Communications Security (CCS '02), Nov. 2002.

Cliff C. Zou et al., "Email Worm Modeling and Defense", In the 13th ACM International Conference on Computer Communications and Networks (CCCN '04), Oct. 2004.

Cliff Changchun Zou et al., "Monitoring and Early Warning for Internet Worms", In Proceedings of the 10th ACM Conference on Computer and Communications Security (CCS '03), Oct. 2003.

Cliff Changchun Zou et al., "On the Performance of Internet Worm Scanning Strategies", Technical Report TR-03-CSE-07, UMass ECE Dept., Nov. 2003.

Zin Zhang et al., "Detecting Stepping Stones", In Proceedings of the 9th USENIX Security Symposium, Aug. 2000.

Alexander Gostev, "Malware Elovution: Jan.-Mar. 2005", Viruslist.com, <http://www.viruslist.com/en/analysis?pubid=162454316>, (Mar. 2005).

Jiang Wu et al., "An Effective Architecture and Algorithm for Detecting Worms with Various Scan Techniques", In Proceedings of the 11th Annual Network and Distributed System Security Symposium (NDSS '04), Feb. 2004.

Matthew M. Williamson et al., "Virus Throttling for Instant Messaging", Virus Bulletin Conference, Sep. 2004, Chicago, IL, USA, (Sep. 2004).

F. Weimer, "Passive DNS Replication", <http://www.enyo.de/fw/software/dnslgger>, 2005.

Ke Wang et al., "Anomalous Payload-Based Network Intrusion Detection", In Proceedings of the 7th International Symposium on Recent Advances in Intrusion Detection (RAID 2004), 2004.

P. Vixie et al., "RFC 2136: Dynamic Updates in the Domain Name System (DNS Update)", <http://www.faqs.org/rfcs/rfc2136.html> (Apr. 1997).

Joe Stewart, "Dipnet/Oddbob Worm Analysis", SecureWorks, <http://www.secureworks.com/research/threats/dipnet/> (Jan. 13, 2005).

Harold Thimbleby et al., "A Framework for Modelling Trojans and Computer Virus Infection", Computer Journal, vol. 41, No. 7, pp. 444-458 (1999).

Paul Bachner et al., "Know Your Enemy: Tracking Botnets", <http://www.honeynet.org/papers/bots/>, (Mar. 13, 2005).

"LockDown Security Bulletin—Sep. 23, 2001", <http://lockdowncorp.com/bots/> (Sep. 23, 2001).

Colleen Shannon et al., "The Spread of the Witty Worm", <http://www.caida.org/analysis/security/witty/index.xml> (Mar. 19, 2004).

Moheeb Abu Rajab et al., "On the Effectiveness of Distributed Worm Monitoring", In Proceedings of the 14th USENIX Security Symposium (2005).

Niels Provos, "CITI Technical Report 03-1: A Virtual Honeypot Framework", <http://www.citi.umich.edu/techreports/reports/citi-tr-03-1.pdf> (Oct. 21, 2003).

"Know your Enemy: Honeynets", <http://www.honeypot.org/papers/honeynet>, (May 31, 2006).

David Moore et al., "Internet Quarantine: Requirements for Containing Self-Propagating Code", In Proceedings of the IEEE INFOCOM 2003, Mar. 2003.

Joe Stewart, "I-Worm Baba Analysis", <http://secureworks.com/research/threats/baba> (Oct. 22, 2004).

David Moore et al., "Slammer Worm Dissection: Inside the Slammer Worm", IEEE Security & Privacy, vol. 1, No. 4 (Jul.-Aug. 2003).

David Moore et al., "Code-Red: A Case Study on the Spread and Victims of an Internet Worm", <http://www.icir.org/vern/imw-2002/imw2002-papers/209.ps.gz> (2002).

Joe Stewart, "Sinit P2P Trojan Analysis", <http://www.secureworks.com/research/threats/sinit>, (Dec. 8, 2003).

Martin Krzywinski, "Port Knocking—Network Authentication Across Closed Ports", Sys. Admin Magazine, vol. 12, pp. 12-17 (2003).

Christopher Kruegel et al., "Anomaly Detection of Web-Based Attacks", In Proceedings of the 10th ACM Conference on Computer and Communication Security (CCS '03), Oct. 27-31, 2003, Washington, DC, USA, pp. 251-261.

"Dabber Worm Analysis", LURHQ Threat Intelligence Group, <http://www.lurhq.com/dabber.html> (May 13, 2004).

Abstract of Jeffrey O. Kephart et al., "Directed-Graph Epidemiological Models of Computer Viruses", Proceedings of the 1991 IEEE Computer Society Symposium on Research in Security and Privacy, Oakland, CA, May 20-22, 1991; pp. 343-359 (May 20-22, 1991).

(56)

References Cited

OTHER PUBLICATIONS

- C. Kalt “RFC 2810—Internet Relay Chat: Architecture” <http://faqs.org/rfcs/rfc2810.html> (Apr. 2000).
- Xuxian Jiang et al., “Cerias Tech Report 2005-24: Virtual Playgrounds for Worm Behavior Investigation”, Purdue University, Feb. 2005.
- Neal Hindocha et al., “Malicious Threats and Vulnerabilities in Instant Messaging”, Virus Bulletin International Conference, Sep. 2003.
- “NSTX (IP-over-DNS) HOWTO”, <http://thomer.com/howtos/nstx.html> (Nov. 4, 2005).
- Christopher W. Hanna, “Using Snort to Detect Rogue IRC Bot Programs”, Technical Report, SANS Institute 2004 (Oct. 8, 2004).
- Jaeyeon Jung et al., “An Empirical Study of Spam Traffic and the Use of DNS Black Lists”, In Proc. ACM SIGCOMM Internet Measurement Conference (IMC '04), Oct. 25-27, 2004, Taormina, Sicily, Italy, pp. 370-375.
- Srikanth Kandula et al., “Botz-4-Sale: Surviving Organized DDoS Attacks That Mimic Flash Crowds”, Technical Report LCS TR-969, Laboratory for Computer Science, MIT, 2004.
- Sven Krasser et al., “Real-Time and Forensic Network Data Analysis Using Animated and Coordinated Visualization”, Proceedings of the 6th IEEE Information Assurance Workshop (Jun. 2005).
- David Moore et al., “Inferring Internet Denial-of-Service Activity”, In Proceedings of the 2001 USENIX Security Symposium, 2001.
- Stephane Racine, “Master’s Thesis: Analysis for Internet Relay Chat Usage by DDoS Zombies”, <ftp://www.tik.ee.ethz.ch/pub/students/2003-2004/Wi/MA-2004-01.pdf> (Nov. 3, 2003).
- Anirudh Ramachandran et al., “Understanding the Network-Level Behavior of Spammers”, SIGCOMM '06, Sep. 11-15, 2006, Pisa, Italy, pp. 291-302.
- Ramneek Puri, “Bots & Botnet: An Overview”, SANS Institute 2003, http://www.giac.com/practical/GSEC/Ramneek_Puri_GSEC.pdf (Aug. 8, 2003).
- Stuart E. Schechter et al., “Access for Sale: A New Class of Worm”, In 2003 ACM Workshop on Rapid Malcode (WORM '03), ACM SIGSAC, Oct. 27, 2003, Washington, DC, USA.
- Stuart Staniford, “How to Own the Internet in Your Spare Time”, In Proc. 11th USENIX Security Symposium, San Francisco, CA, Aug. 2002.
- Martin Overton, “Bots and Botnets: Risks, Issues and Prevention”, 2005 Virus Bulletin Conference at the Burlington, Dublin, Ireland, Oct. 5-7, 2005, http://arachnid.homeip.net/papers/VB2005-Bots_and_Botnets-1.0.2.pdf.
- Yin Zhang et al., “Detecting Stepping Stones”, Proceedings of the 9th USENIX Security Symposium, Denver, Colorado, USA, Aug. 14-17, 2000.
- David Dagon et al., “Worm Population Control Through Periodic Response”, Technical Report, Georgia Institute for Technology, Jun. 2004.
- Scott Jones et al., “The IPM Model of Computer Virus Management”, Computers & Security, vol. 9, pp. 411-418 (1990).
- Jeffrey O. Kephart et al., “Directed-Graph Epidemiological Models of Computer Viruses”, In Proceedings of IEEE Symposium on Security and Privacy, pp. 343-359 (1991).
- Darrell M. Kienzle et al., “Recent Worms: A Survey and Trends”, In WORM '03, Proceedings of the 2003 ACM Workshop on Rapid Malcode, Washington, DC, USA, pp. 1-10, Oct. 27, 2003.
- Bill McCarty, “Botnets: Big and Bigger”, IEEE Security and Privacy Magazine, vol. 1, pp. 87-89 (2003).
- Xinzhou Qin et al., “Worm Detection Using Local Networks”, Technical Report GIT-CC-04-04, College of Computing, Georgia Institute of Technology, Feb. 2004.
- Yang Wang et al., “Modeling the Effects of Timing Parameters on Virus Propagation”, In Proceedings of ACM CCS Workshop on Rapid Malcode (WORM '03), Washington, DC, pp. 61-66, Oct. 27, 2003.
- Donald J. Welch et al., “Strike Back: Offensive Actions in Information Warfare”, in AMC New Security Paradigm Workshop, pp. 47-52 (1999).
- T. Liston, “Welcome to My Tarpit: The Tactical and Strategic Use of LaBrea”, <http://www.hackbusters.net/LaBrea/LaBrea.text>, Oct. 24, 2001.
- R. Pointer, “Eggdrop Development”, <http://www.eggheads.org>, Oct. 1, 2005.
- S. Staniford, Code Red Analysis Pages: July Infestation Analysis, <http://www.silicondefense.com/cr/july.html>, Nov. 18, 2001.
- Alex Ma, “NetGeo—The Internet Geographic Database”, <http://www.caida.org/tools/utilities/netgeo/index.xml>, Sep. 6, 2006.
- MathWorks Inc. Simulink, <http://www.mathworks.com/products/simulink>, Dec. 31, 2005.
- David Dagon et al., “Modeling Botnet Propagation Using Time Zones”, In Proceedings of the 13th Annual Network and Distributed Systems Security Symposium (NDSS '06), Feb. 2006.
- John Canavan, “Symantec Security Response: W32.Bobax.D”, <http://www.sarc.com/avcent/venc/data/w32.bobax.d.html>, May 26, 2004.
- “Whois Privacy”, www.gnso.icann.org/issues/whois-privacy/index.shtml, Jun. 3, 2005.
- John D. Hardin, “The Scanner Tarpit HOWTO”, <http://www.impsec.org/linus/security/scanner-tarpit.html>, Jul. 20, 2002.
- Charles J. Krebs, “Ecological Methodology”, Harper & Row, Publishers, New York, pp. v-x, 15-37, 155-166, and 190-194 (1989).
- D.J. Daley et al., “Epidemic Modelling: An Introduction”, Cambridge University Press, pp. vii-ix, 7-15, and 27-38 (1999).
- Lance Spitzner, “Honeypots: Tracking Hackers”, Addison-Wesley, pp. vii-xiv, 73-139, 141-166, and 229-276 (2003).
- V. Fuller et al., “RFC 1519—Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy”, <http://www.faqs.org/rfcs/rfc1519.html> (Sep. 1993).
- David E. Smith “Dynamic DNS”, <http://www.technopagan.org/dynamic> (Aug. 7, 2006).
- Dave Dittrich, “Active Response Continuum Research Project”, <http://staff.washington.edu/dittrich/arc/> (Nov. 14, 2005).
- Joe Stewart, “Akak Trojan Analysis”, <http://www.secureworks.com/research/threats/akak/> (Aug. 31, 2004).
- Monirul I. Sharif, “Mechanisms of Dynamic Analysis and DSTRACE”.
- Kapil Kumar Singh, “IRC Reconnaissance (IRCRecon) Public IRC Heuristics (BotSniffer)” (Jul. 24, 2006). <http://www.trendmicro.com/en/home/us/home.htm>.
- “InterCloud Security Service”, <http://www.trendmicro.com/en/products/nss/icss/evaluate/overview.thm>.
- “2006 Press Releases: Trend Micro Takes Unprecedented Approach to Eliminating Botnet Threats with the Unveiling of InterCloud Security Service”, <http://www.trendmicro.com/en/about/news/pr/archive/2006/pr092506.htm>, (Sep. 25, 2006).
- Paul F. Roberts, “Trend Micro Launches Anti-Botnet Service”, InfoWorld, http://www.infoworld.com/article/06/09/25/HNtrendintercloud_1.html (Sep. 25, 2006).
- CNN Technology News—Expert: Botnets No. 1 Emerging Internet Threat, CNN.com, <http://www.cnn.com/2006/TECH/internet/01/31/furst.index.html> (Jan. 31, 2006).
- Evan Cooke et al., “The Zombie Roundup: Understanding, Detecting, and Disrupting Botnets”, In USENIX Workshop on Steps to Reducing Unwanted Traffic on the Internet (SRUTI), Jun. 2005.
- Sven Dietrich et al., “Analyzing Distributed Denial of Service Tools: The Shaft Case”, Proceedings of the 14th Systems Administration Conference (LISA 2000), New Orleans, Louisiana, USA, Dec. 3-8, 2000.
- Felix C. Freiling et al., “Botnet Tracking: Exploring a Root-Cause Methodology to Prevent Distributed Denial-of-Service Attacks”, ESORICS 2005, LNCS 3679, pp. 319-335 (2005).
- Luiz Henrique Gomes et al., “Characterizing a Spam Traffic”, In Proc. ACM SIGCOMM Internet Measurement Conference (IMC '04), Oct. 25-27, 2004 Taormina, Sicily, Italy, pp. 356-369.
- File History of U.S. Appl. No. 12/985,140, electronically captured on Sep. 3, 2013 for Jun. 3, 2013 to Sep. 3, 2013.
- File History of U.S. Appl. No. 13/205,928, electronically captured on Sep. 3, 2013 for Jun. 3, 2013 to Sep. 3, 2013.
- File History of U.S. Appl. No. 13/358,303, electronically captured on Sep. 3, 2013 for Jun. 3, 2013 to Sep. 3, 2013.

(56)

References Cited

OTHER PUBLICATIONS

<https://sie.isc.org/>, retrieved from Internet Archive on May 23, 2013, Archived Dec. 29, 2008 (2 pages).

“Troj/Agobot-IB”, <http://www.sophos.com/virusinfo/analyses/trojagobotib.html>, printed Jun. 27, 2012 (1 page).

“Norman Sandbox Whitepaper”, Copyright Norman 2003 (19 pages).

<http://www.mcafee.com/us/>, printed May 23, 2013 (23 pages).

“Windows Virtual PC”, http://en.wikipedia.org/wiki/Windows_Virtual_PC, Last Modified May 5, 2013, Printed May 23, 2013 (21 pages).

<http://handlers.sans.org/jclausing/userdb.txt>, printed May 24, 2013 (149 pages).

Timo Sirainen, “IRSSI”, <http://en.wikipedia.org/wiki/Irssi>, updated May 8, 2013 (3 pages).

* cited by examiner

FIGURE 1A

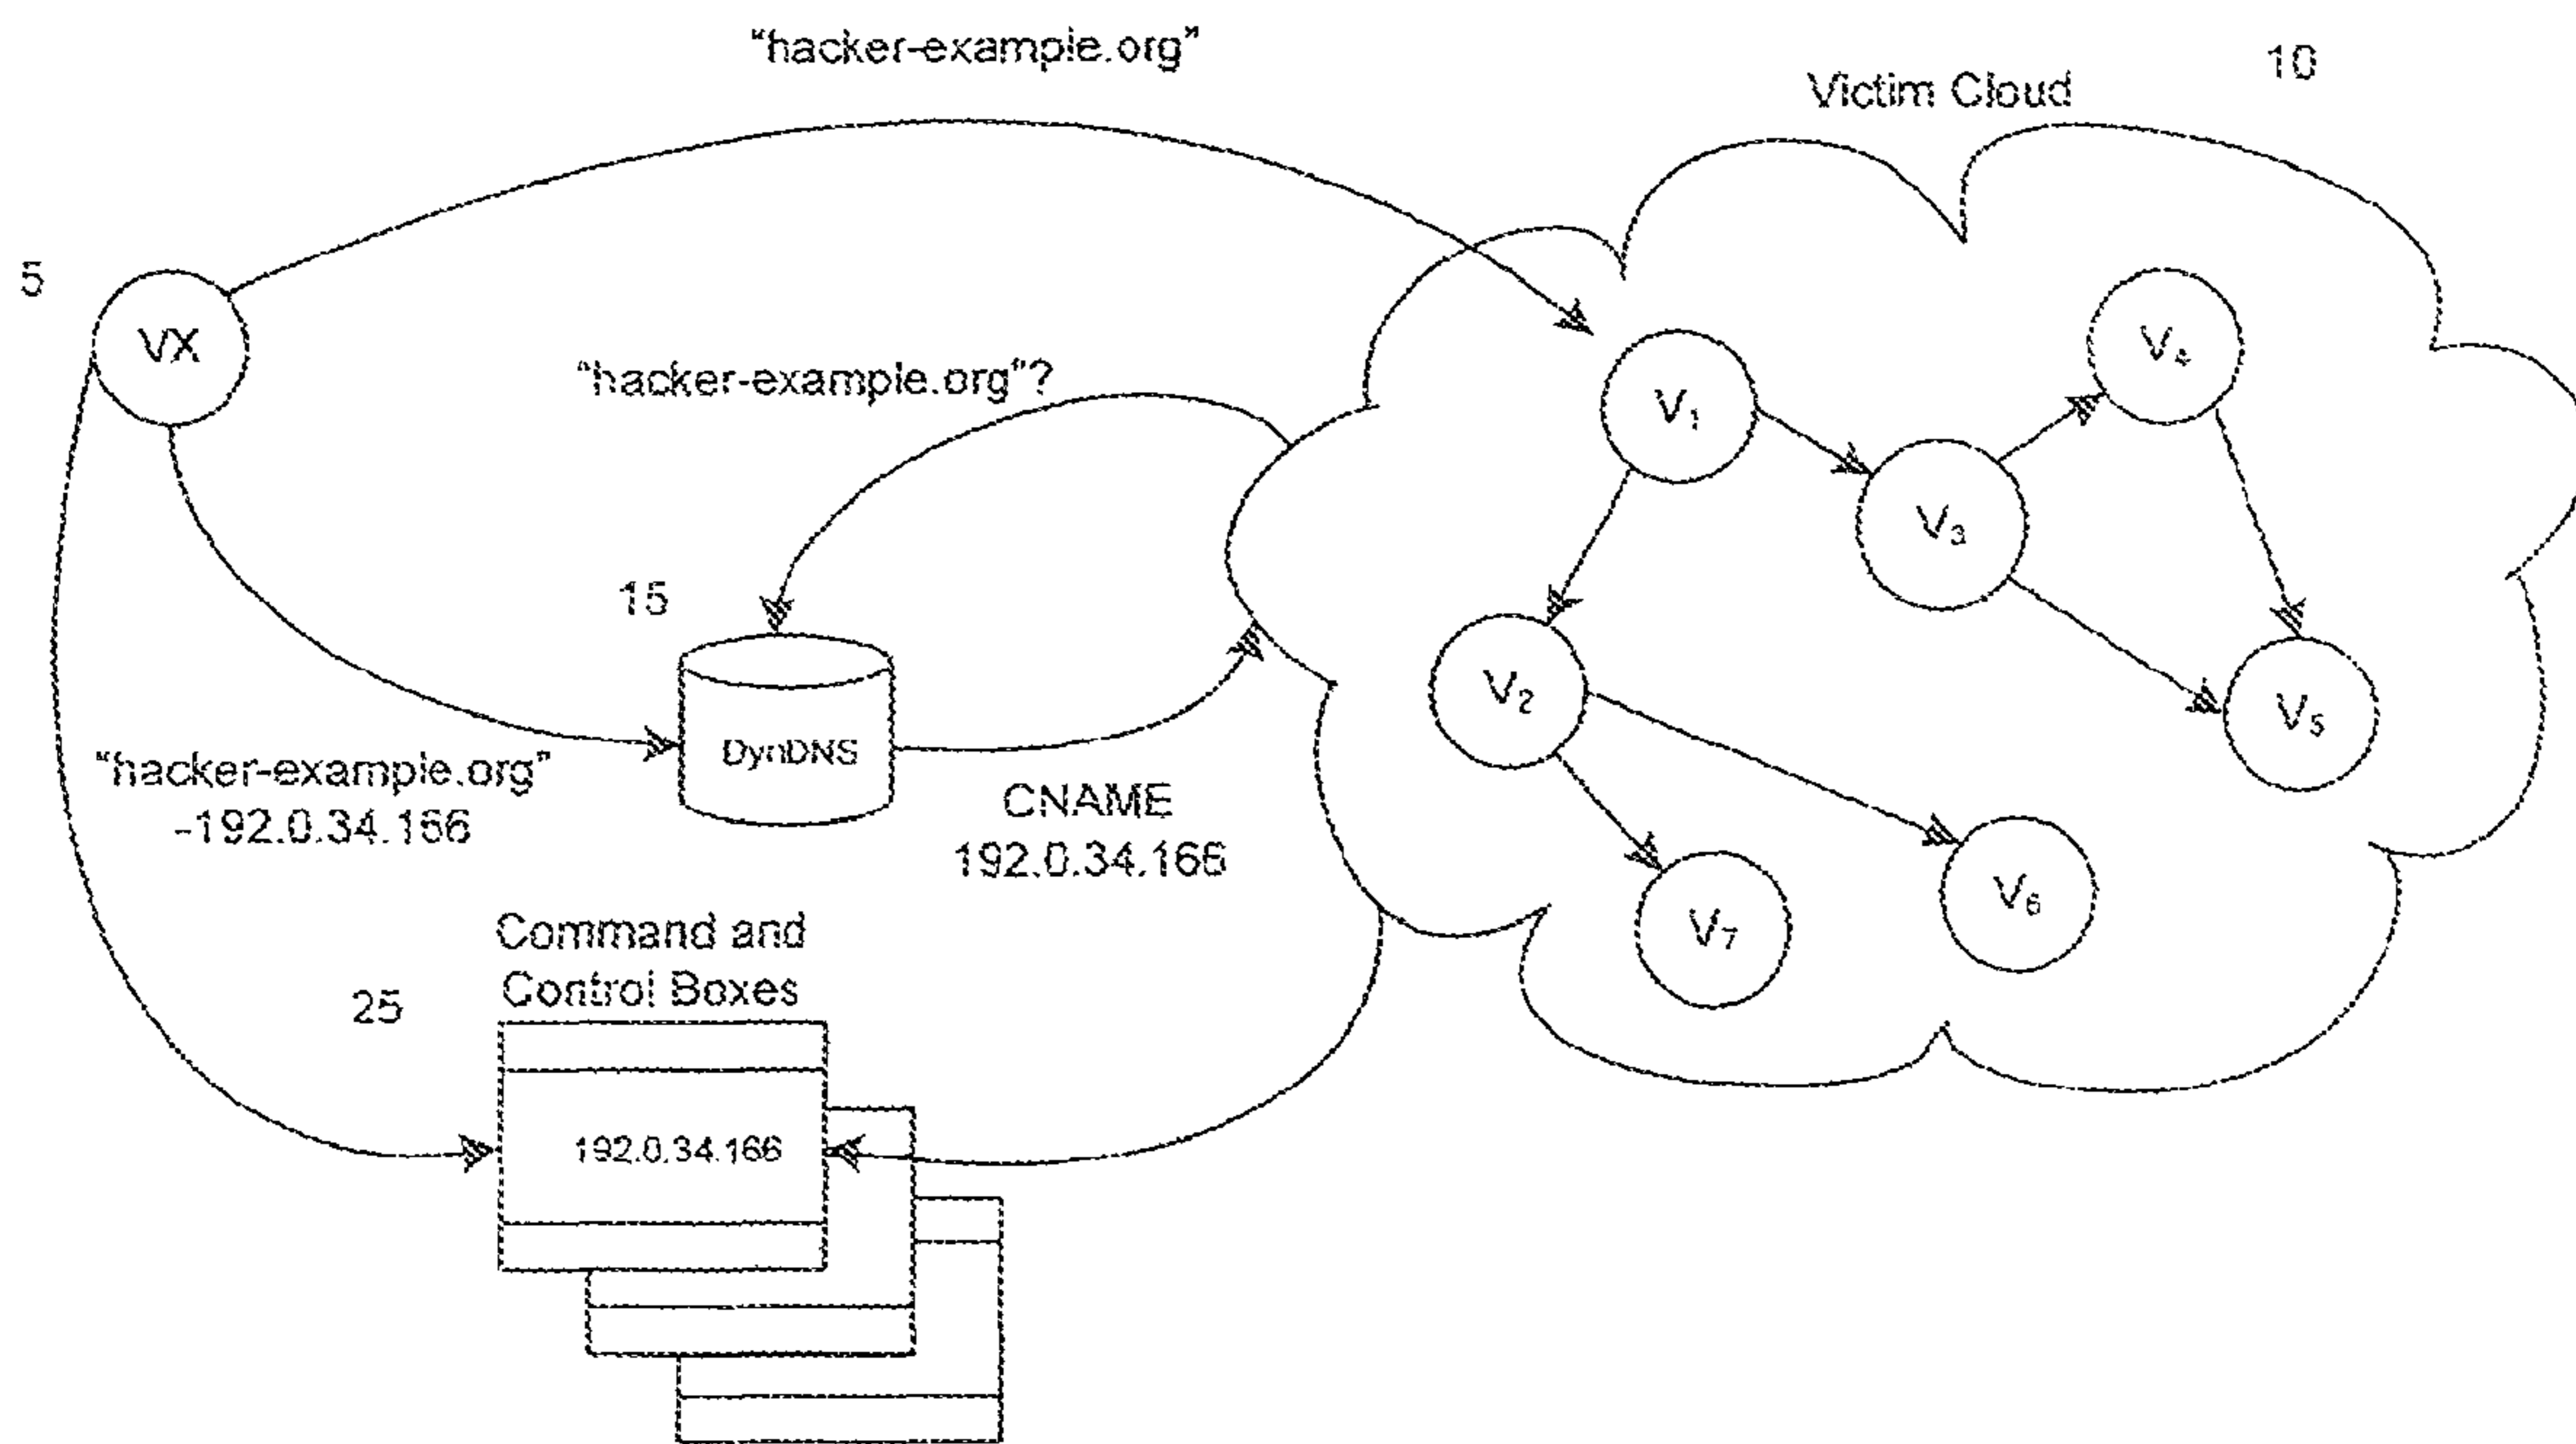


FIGURE 1B

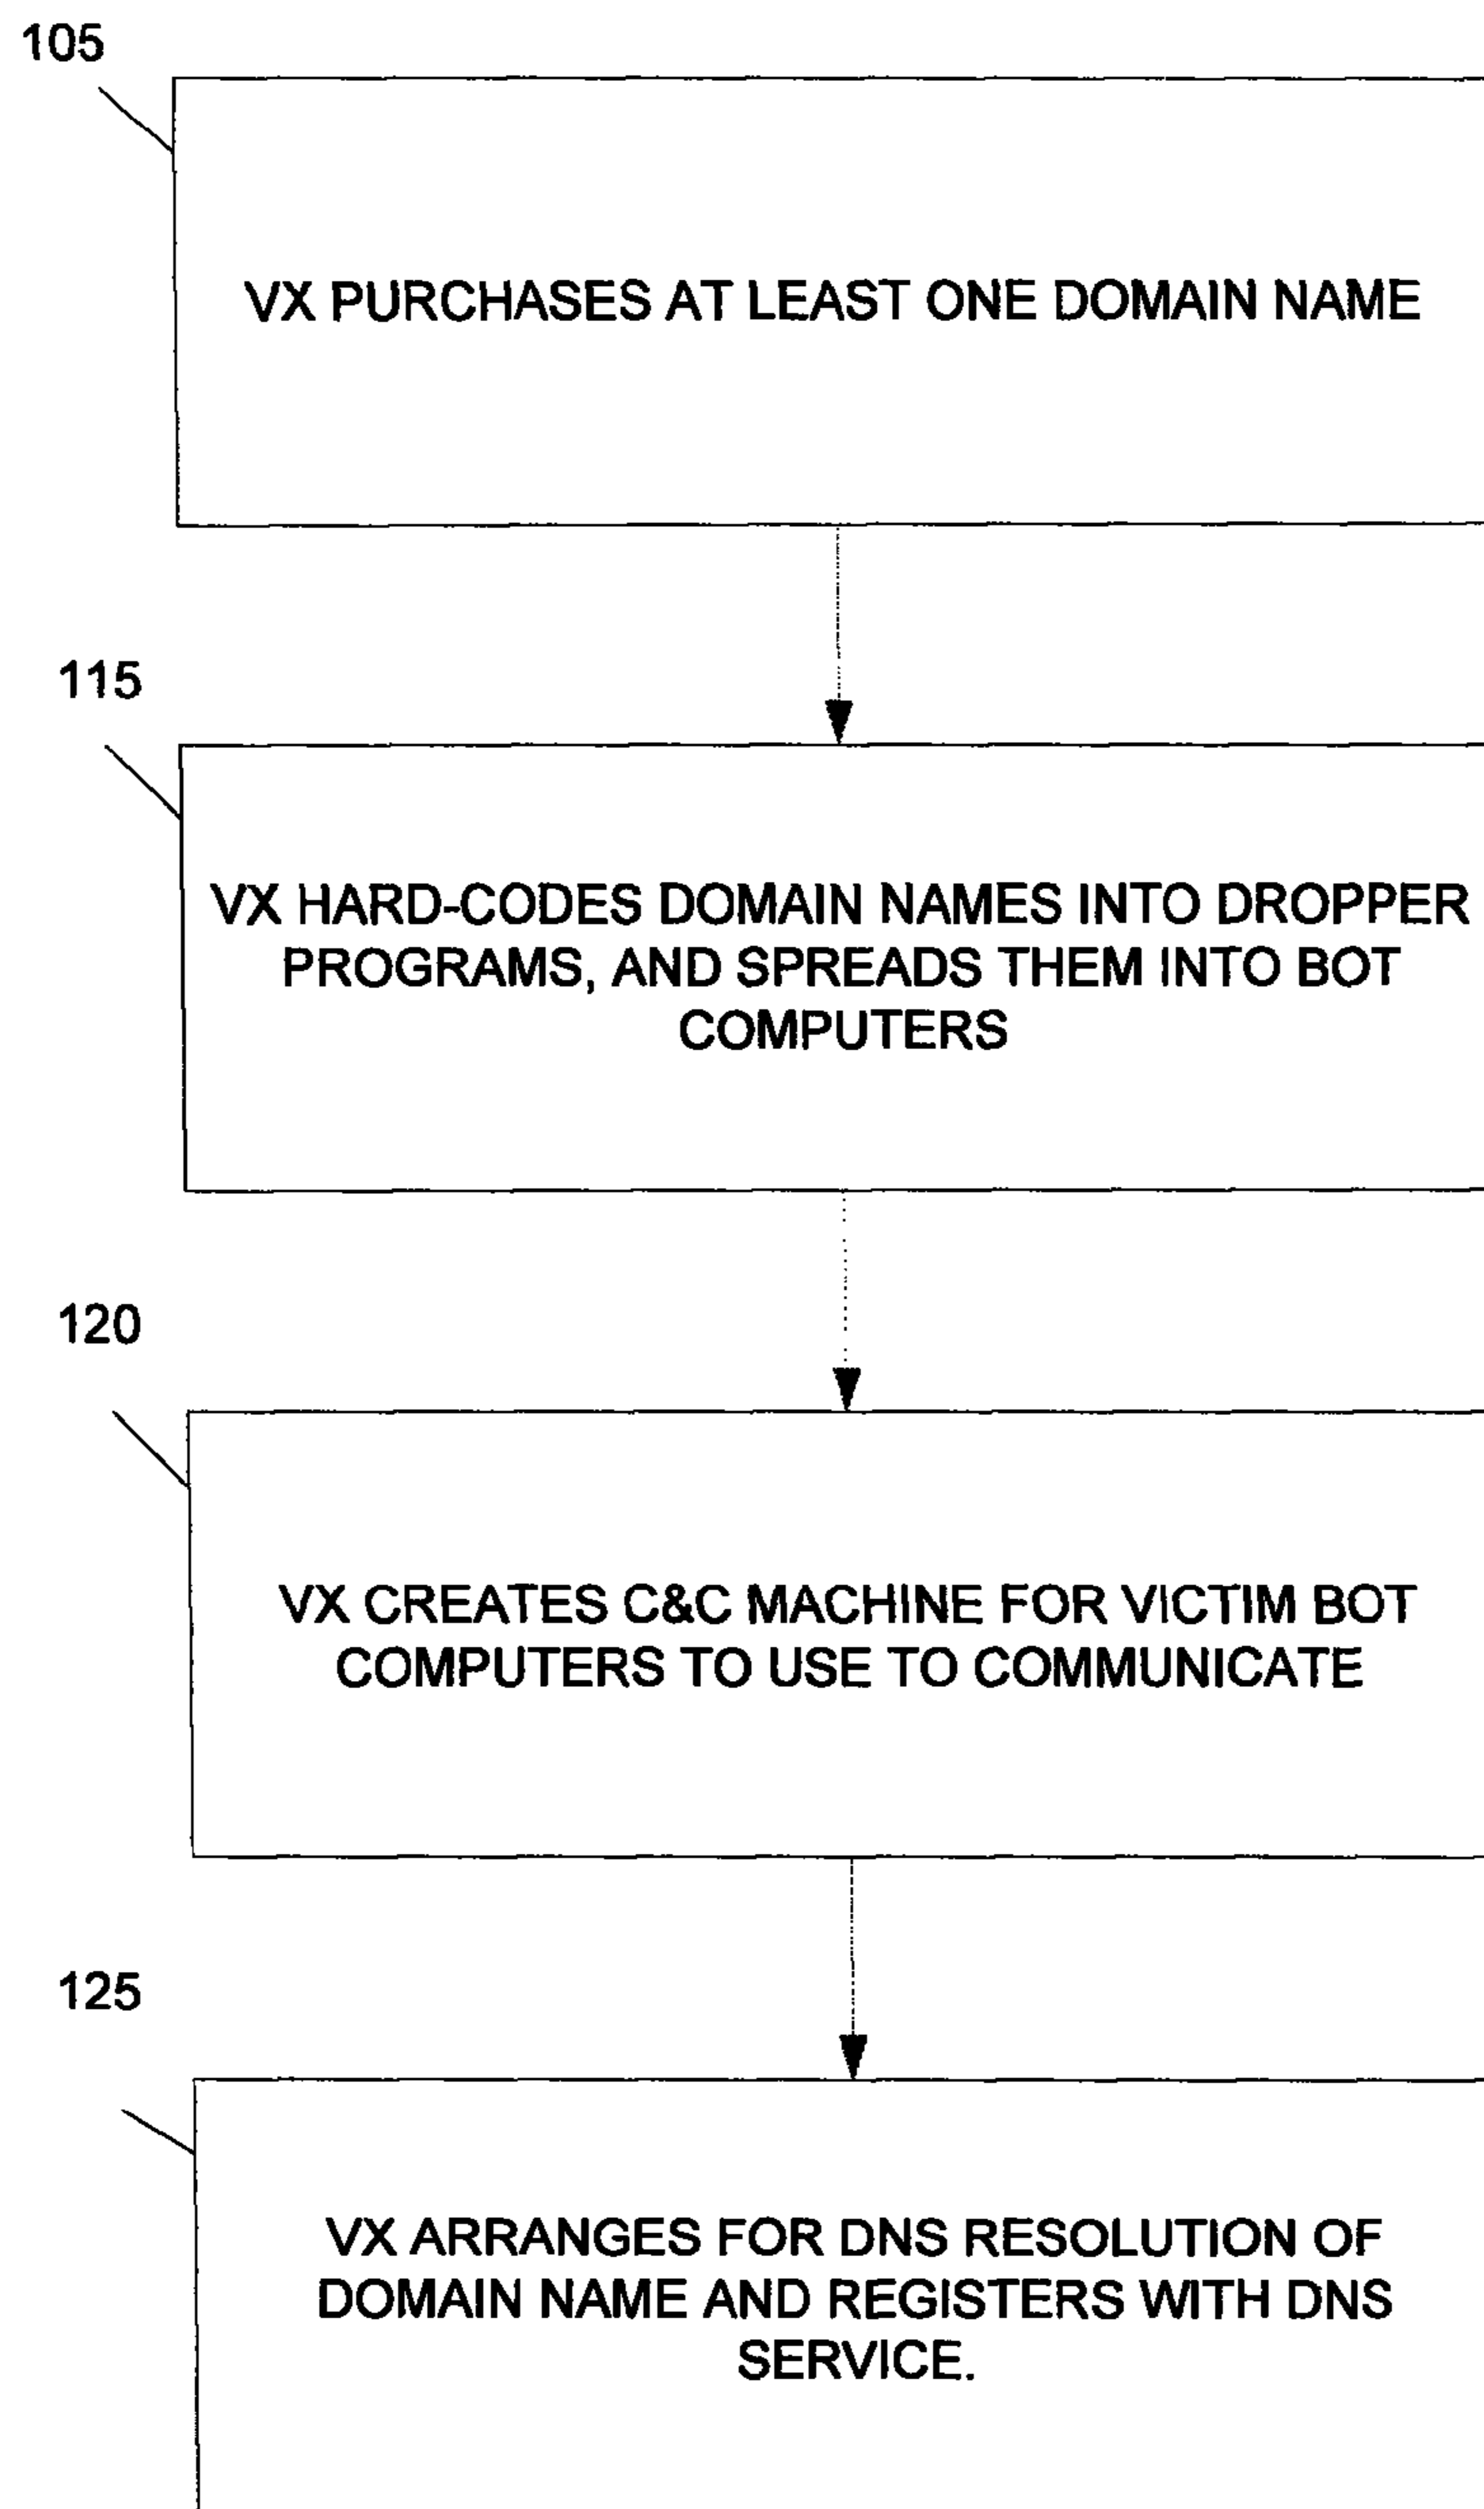


FIGURE 2A

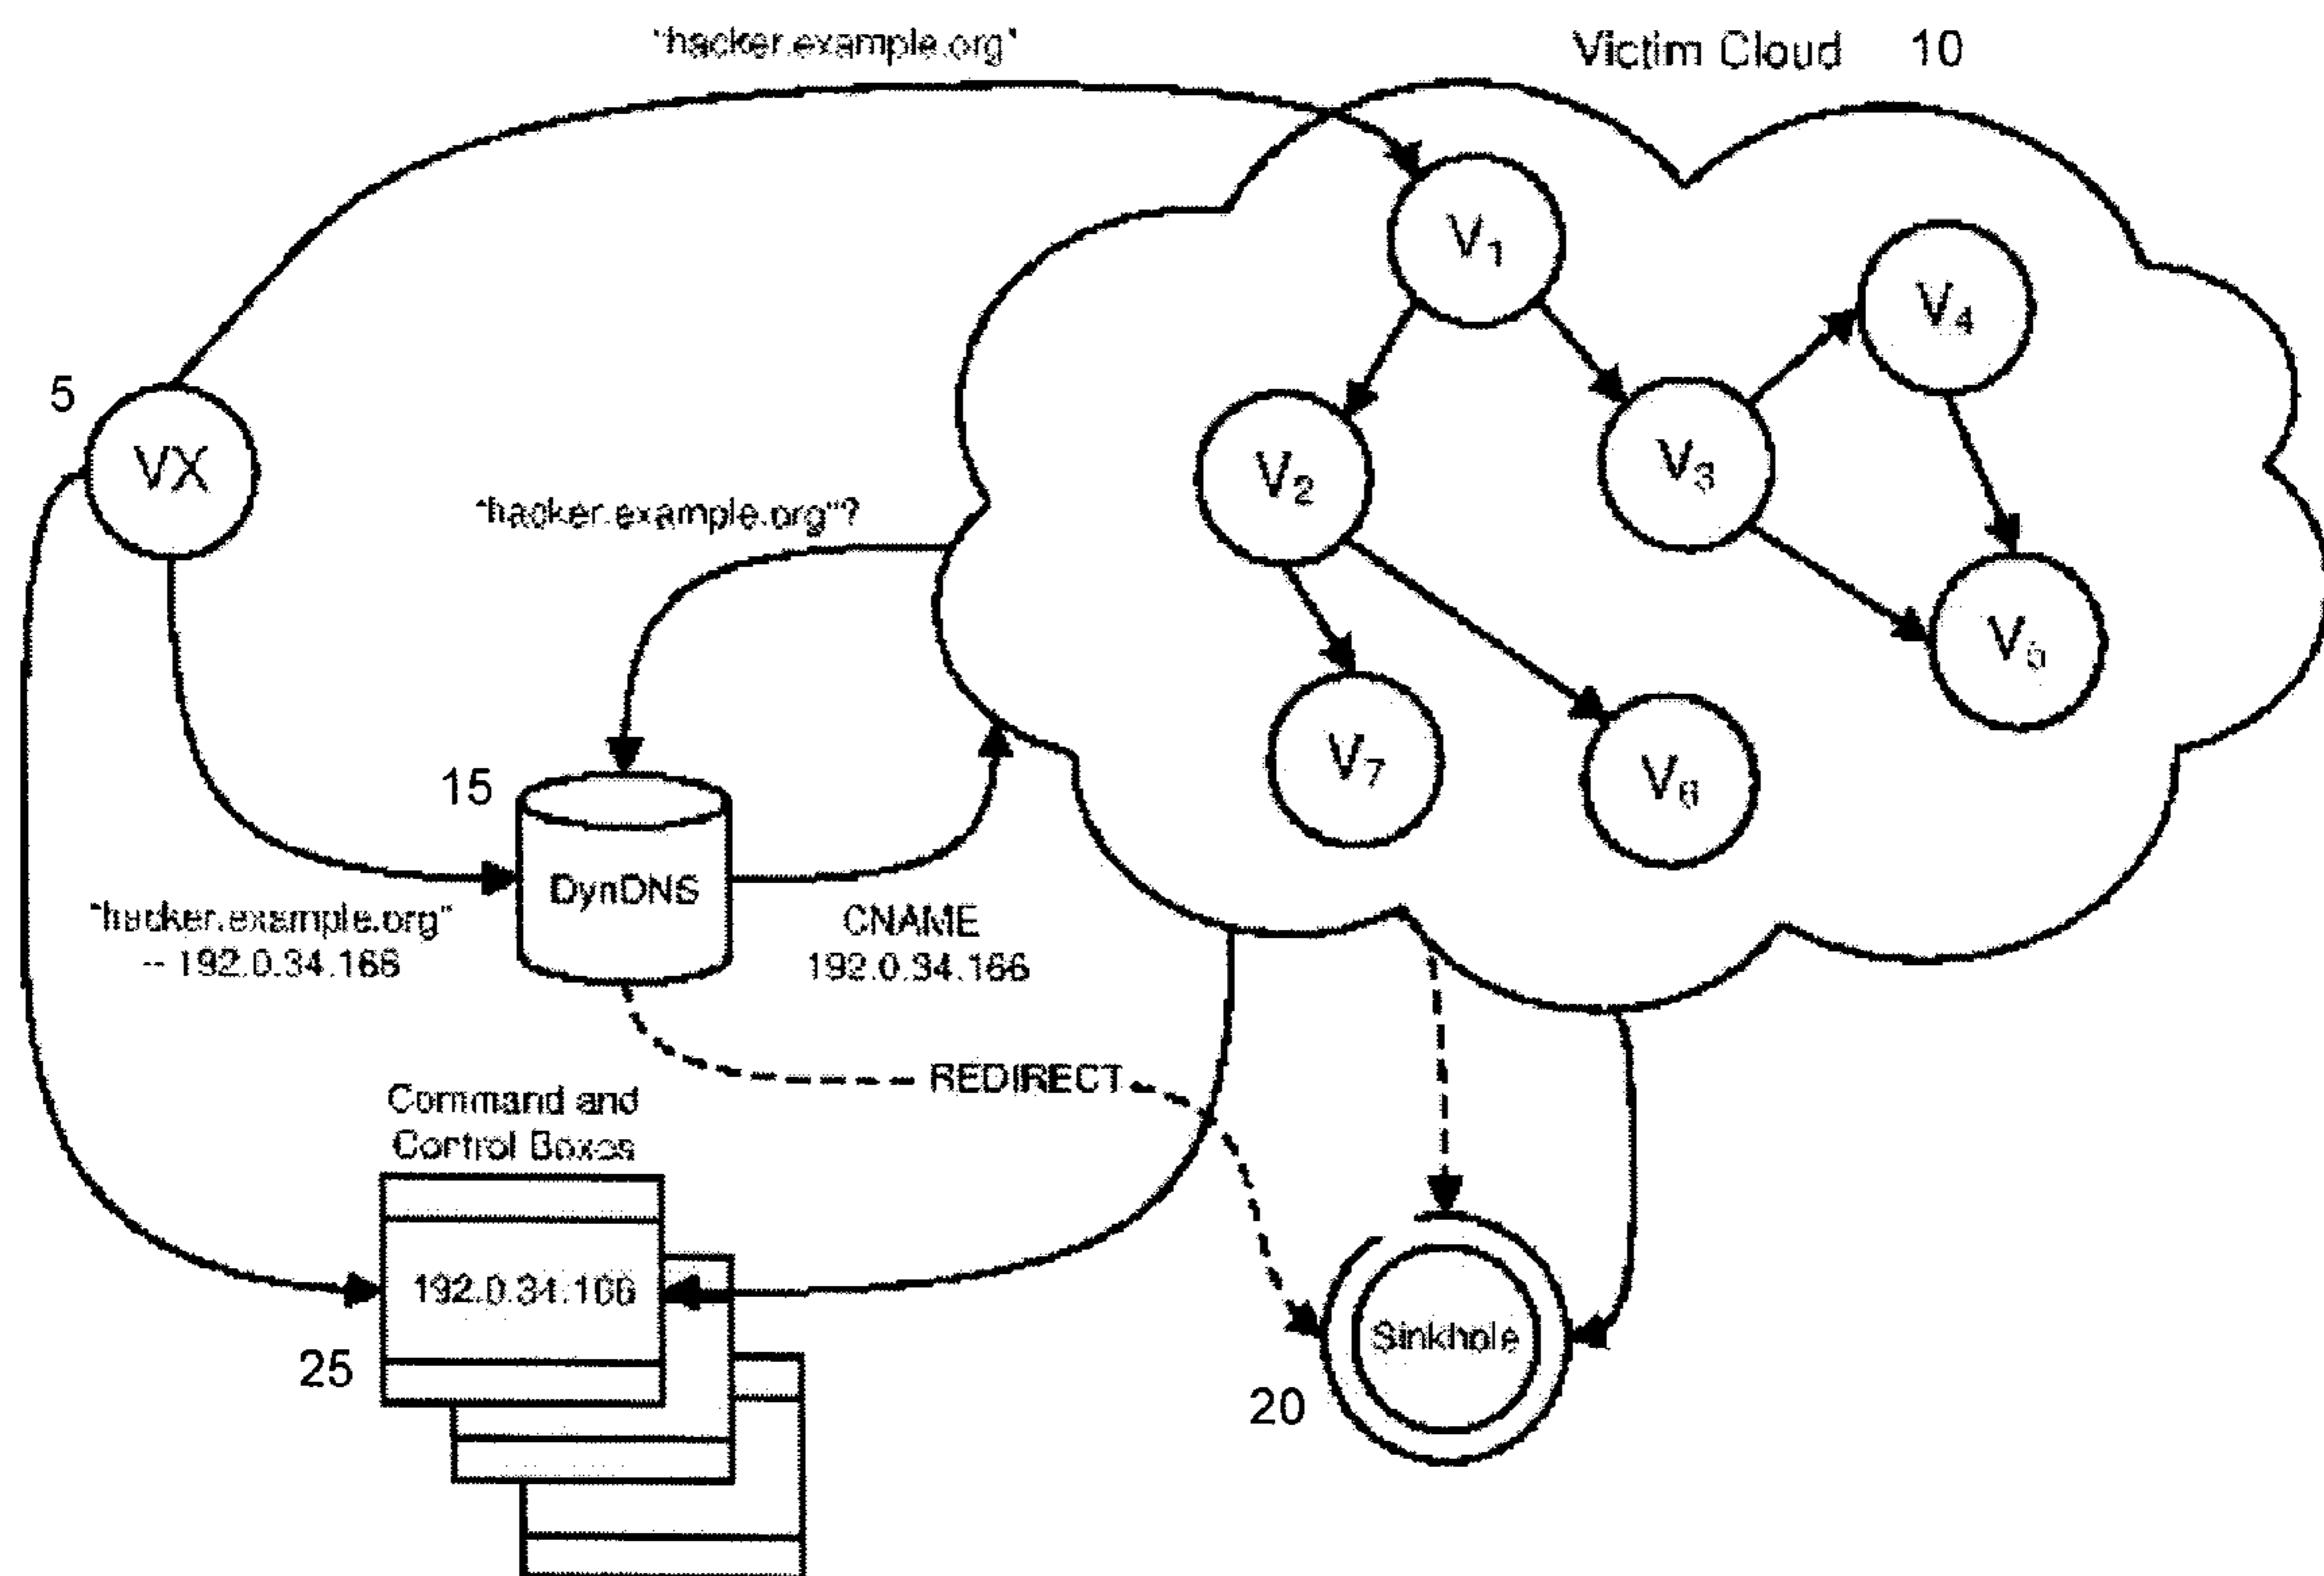


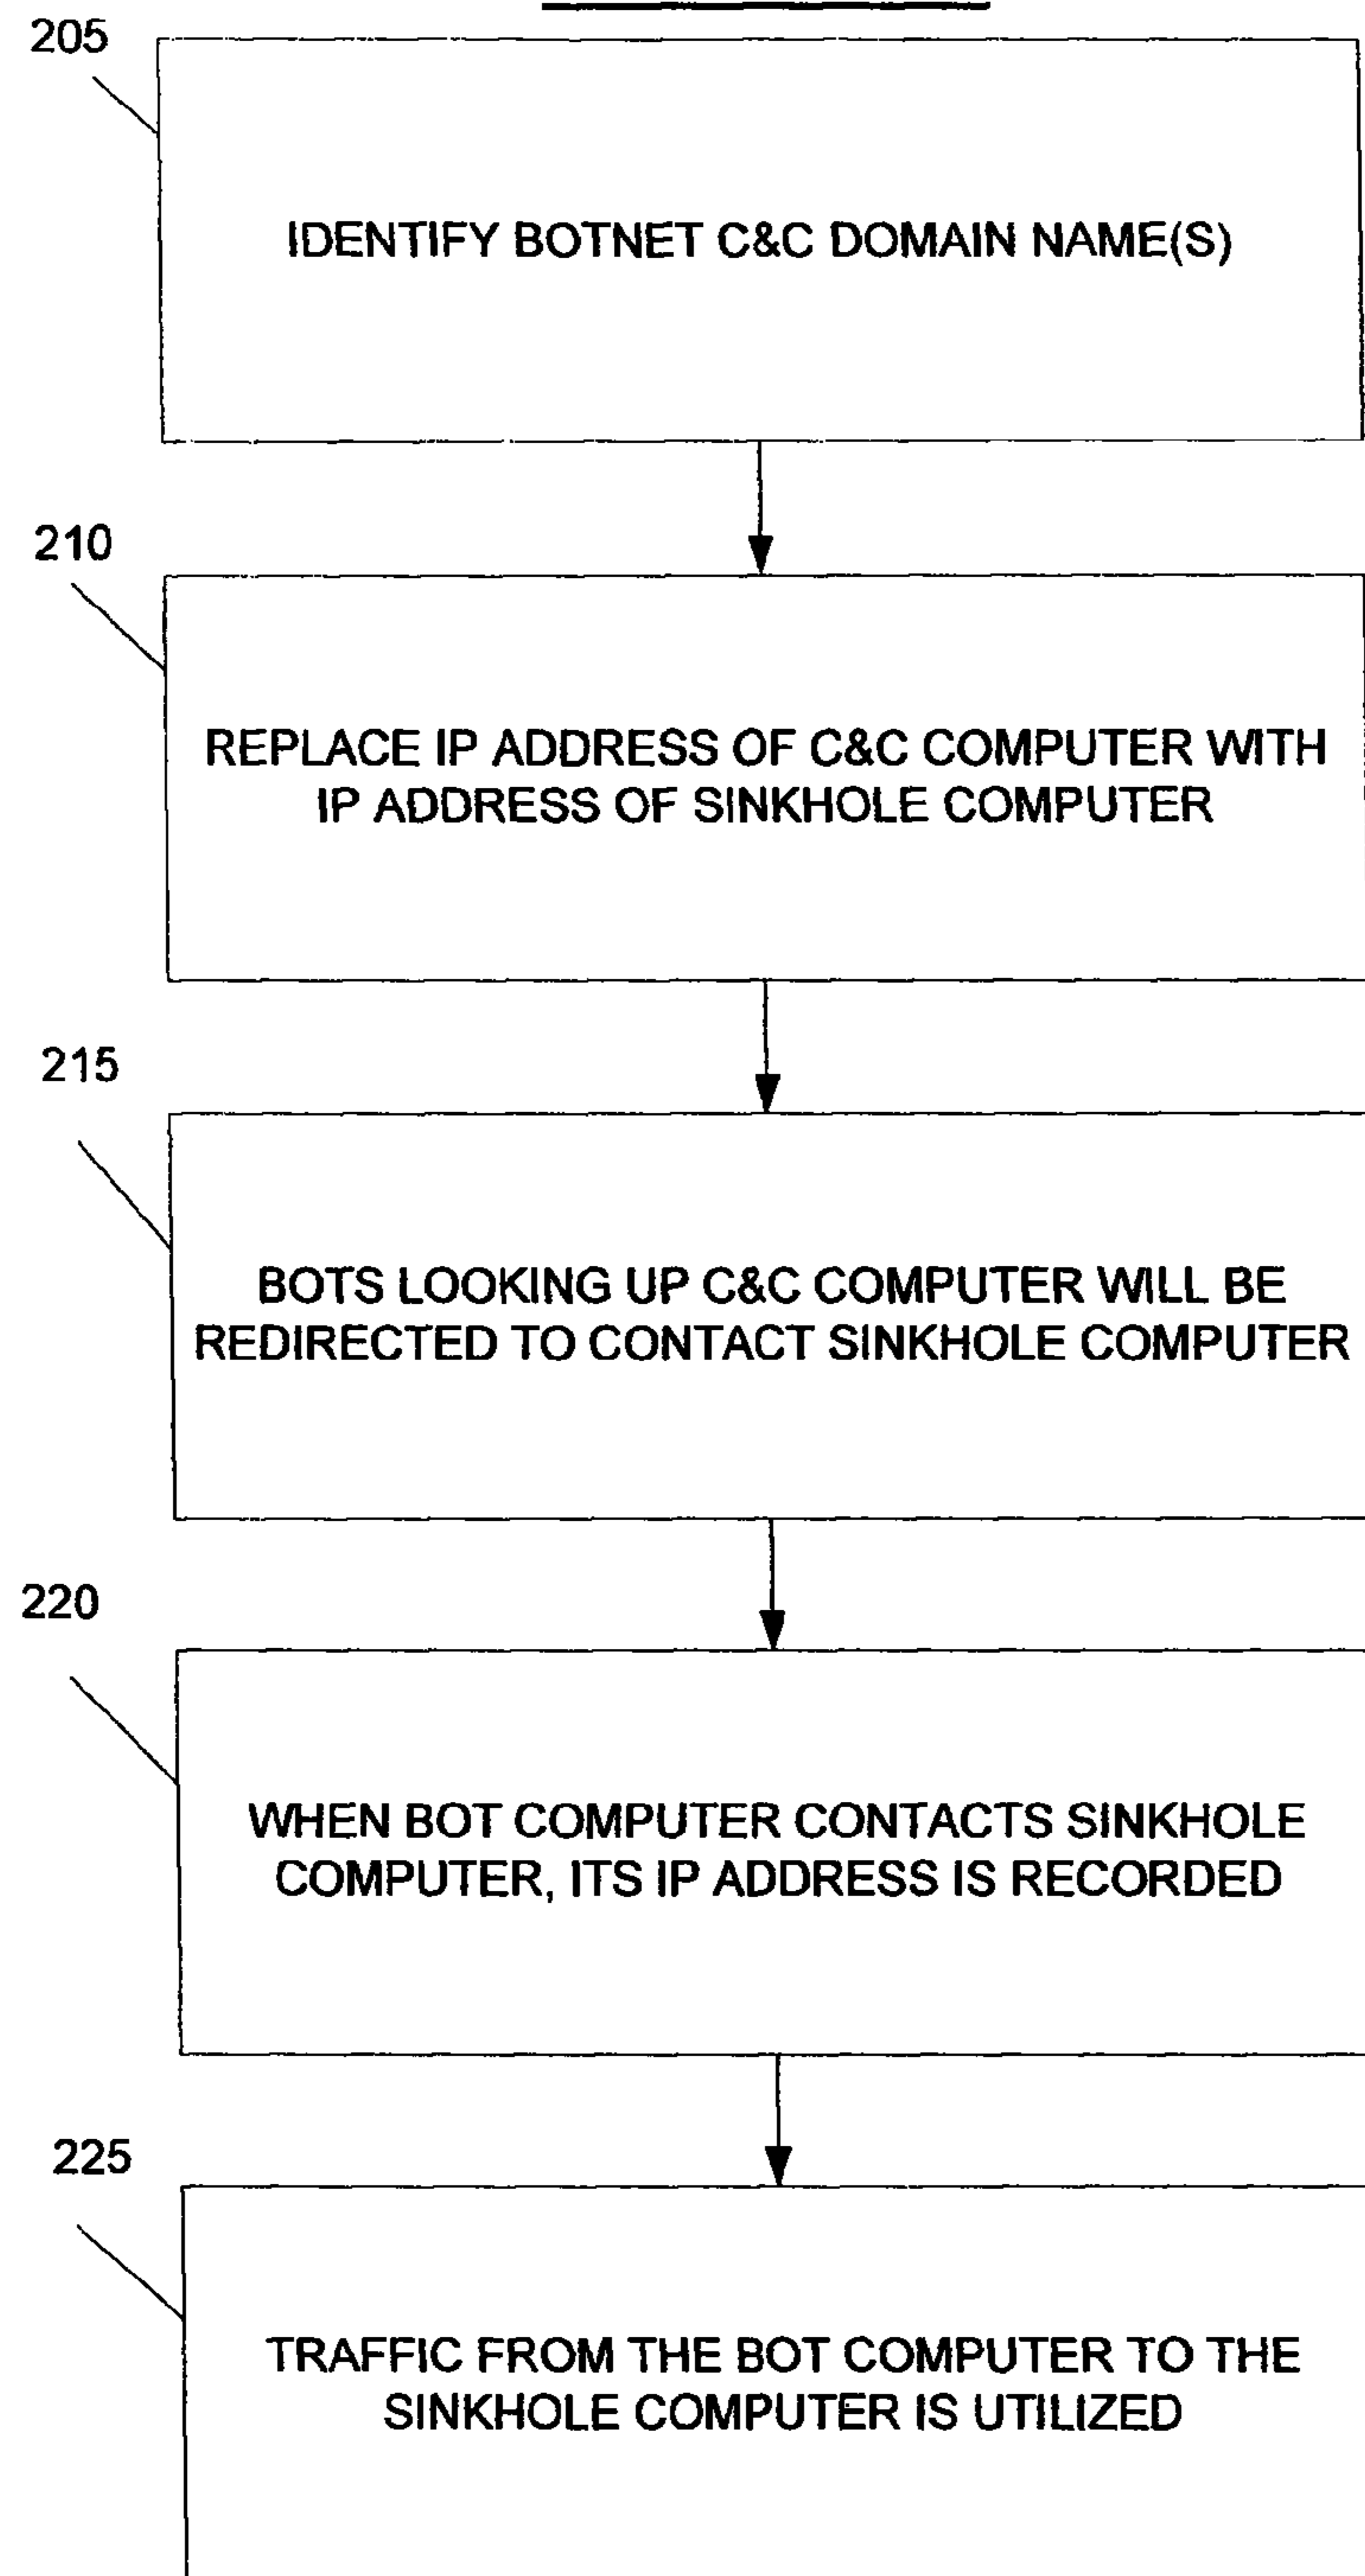
FIGURE 2B

FIGURE 2C

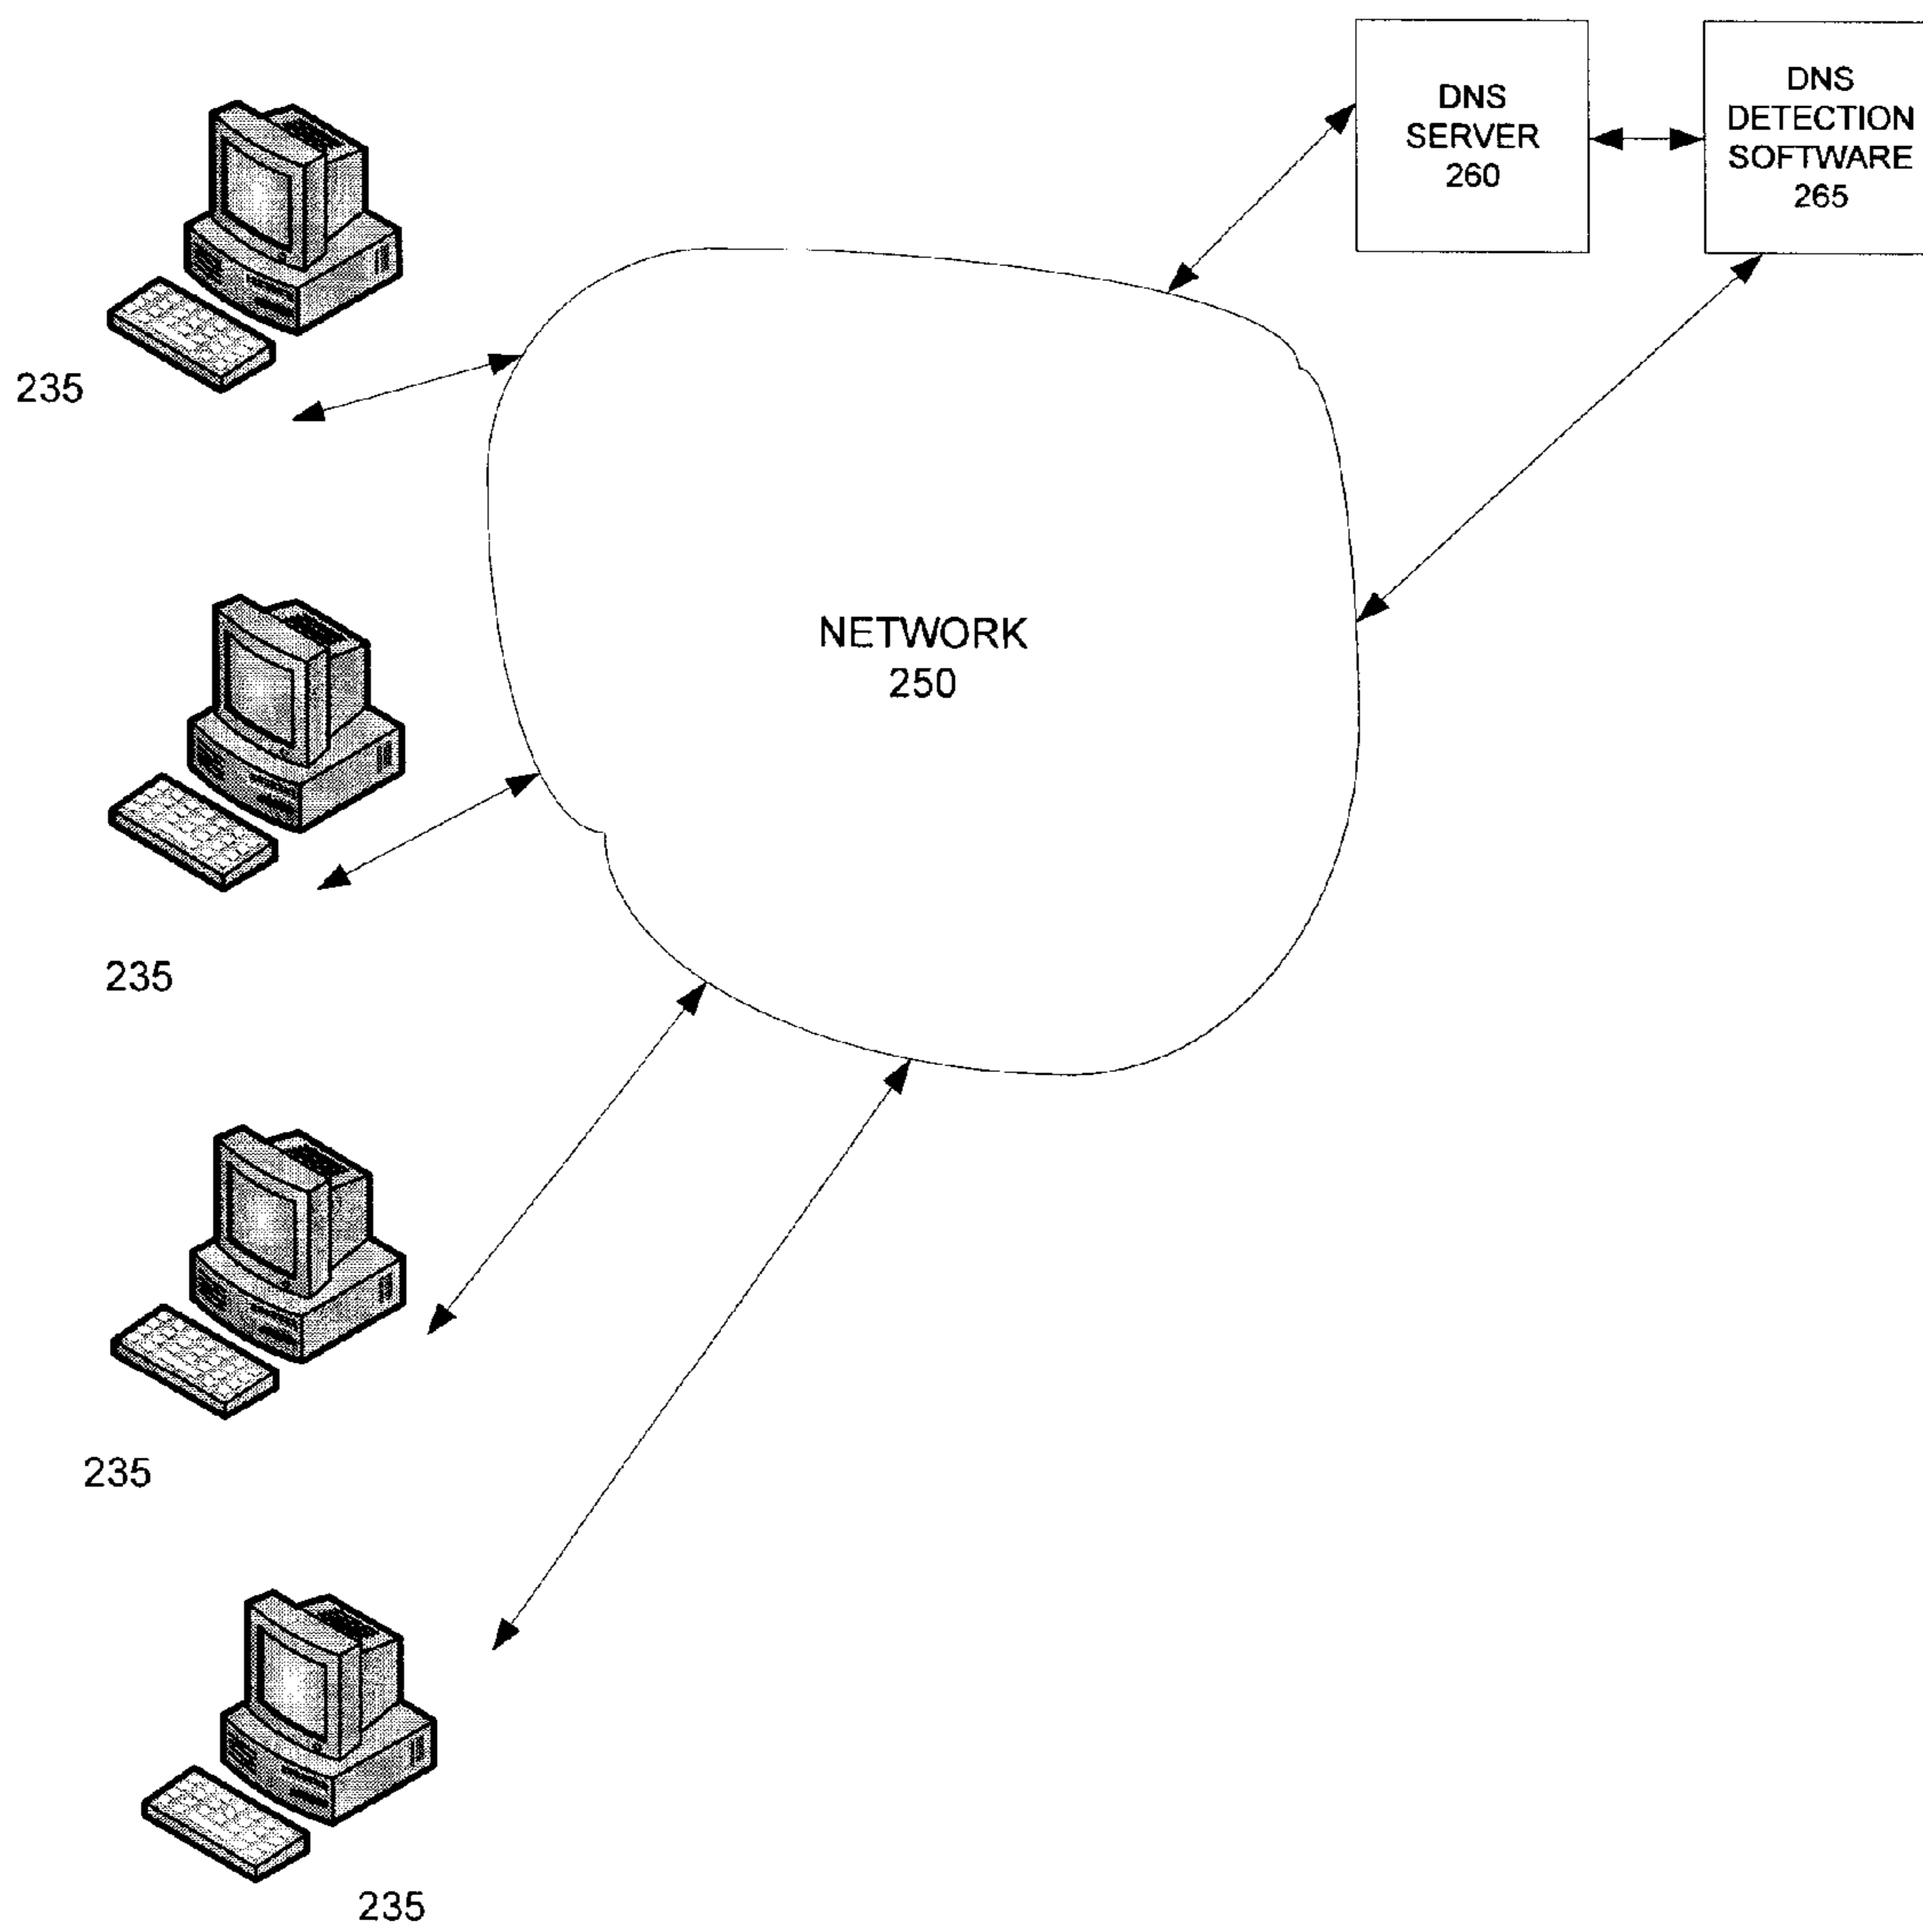
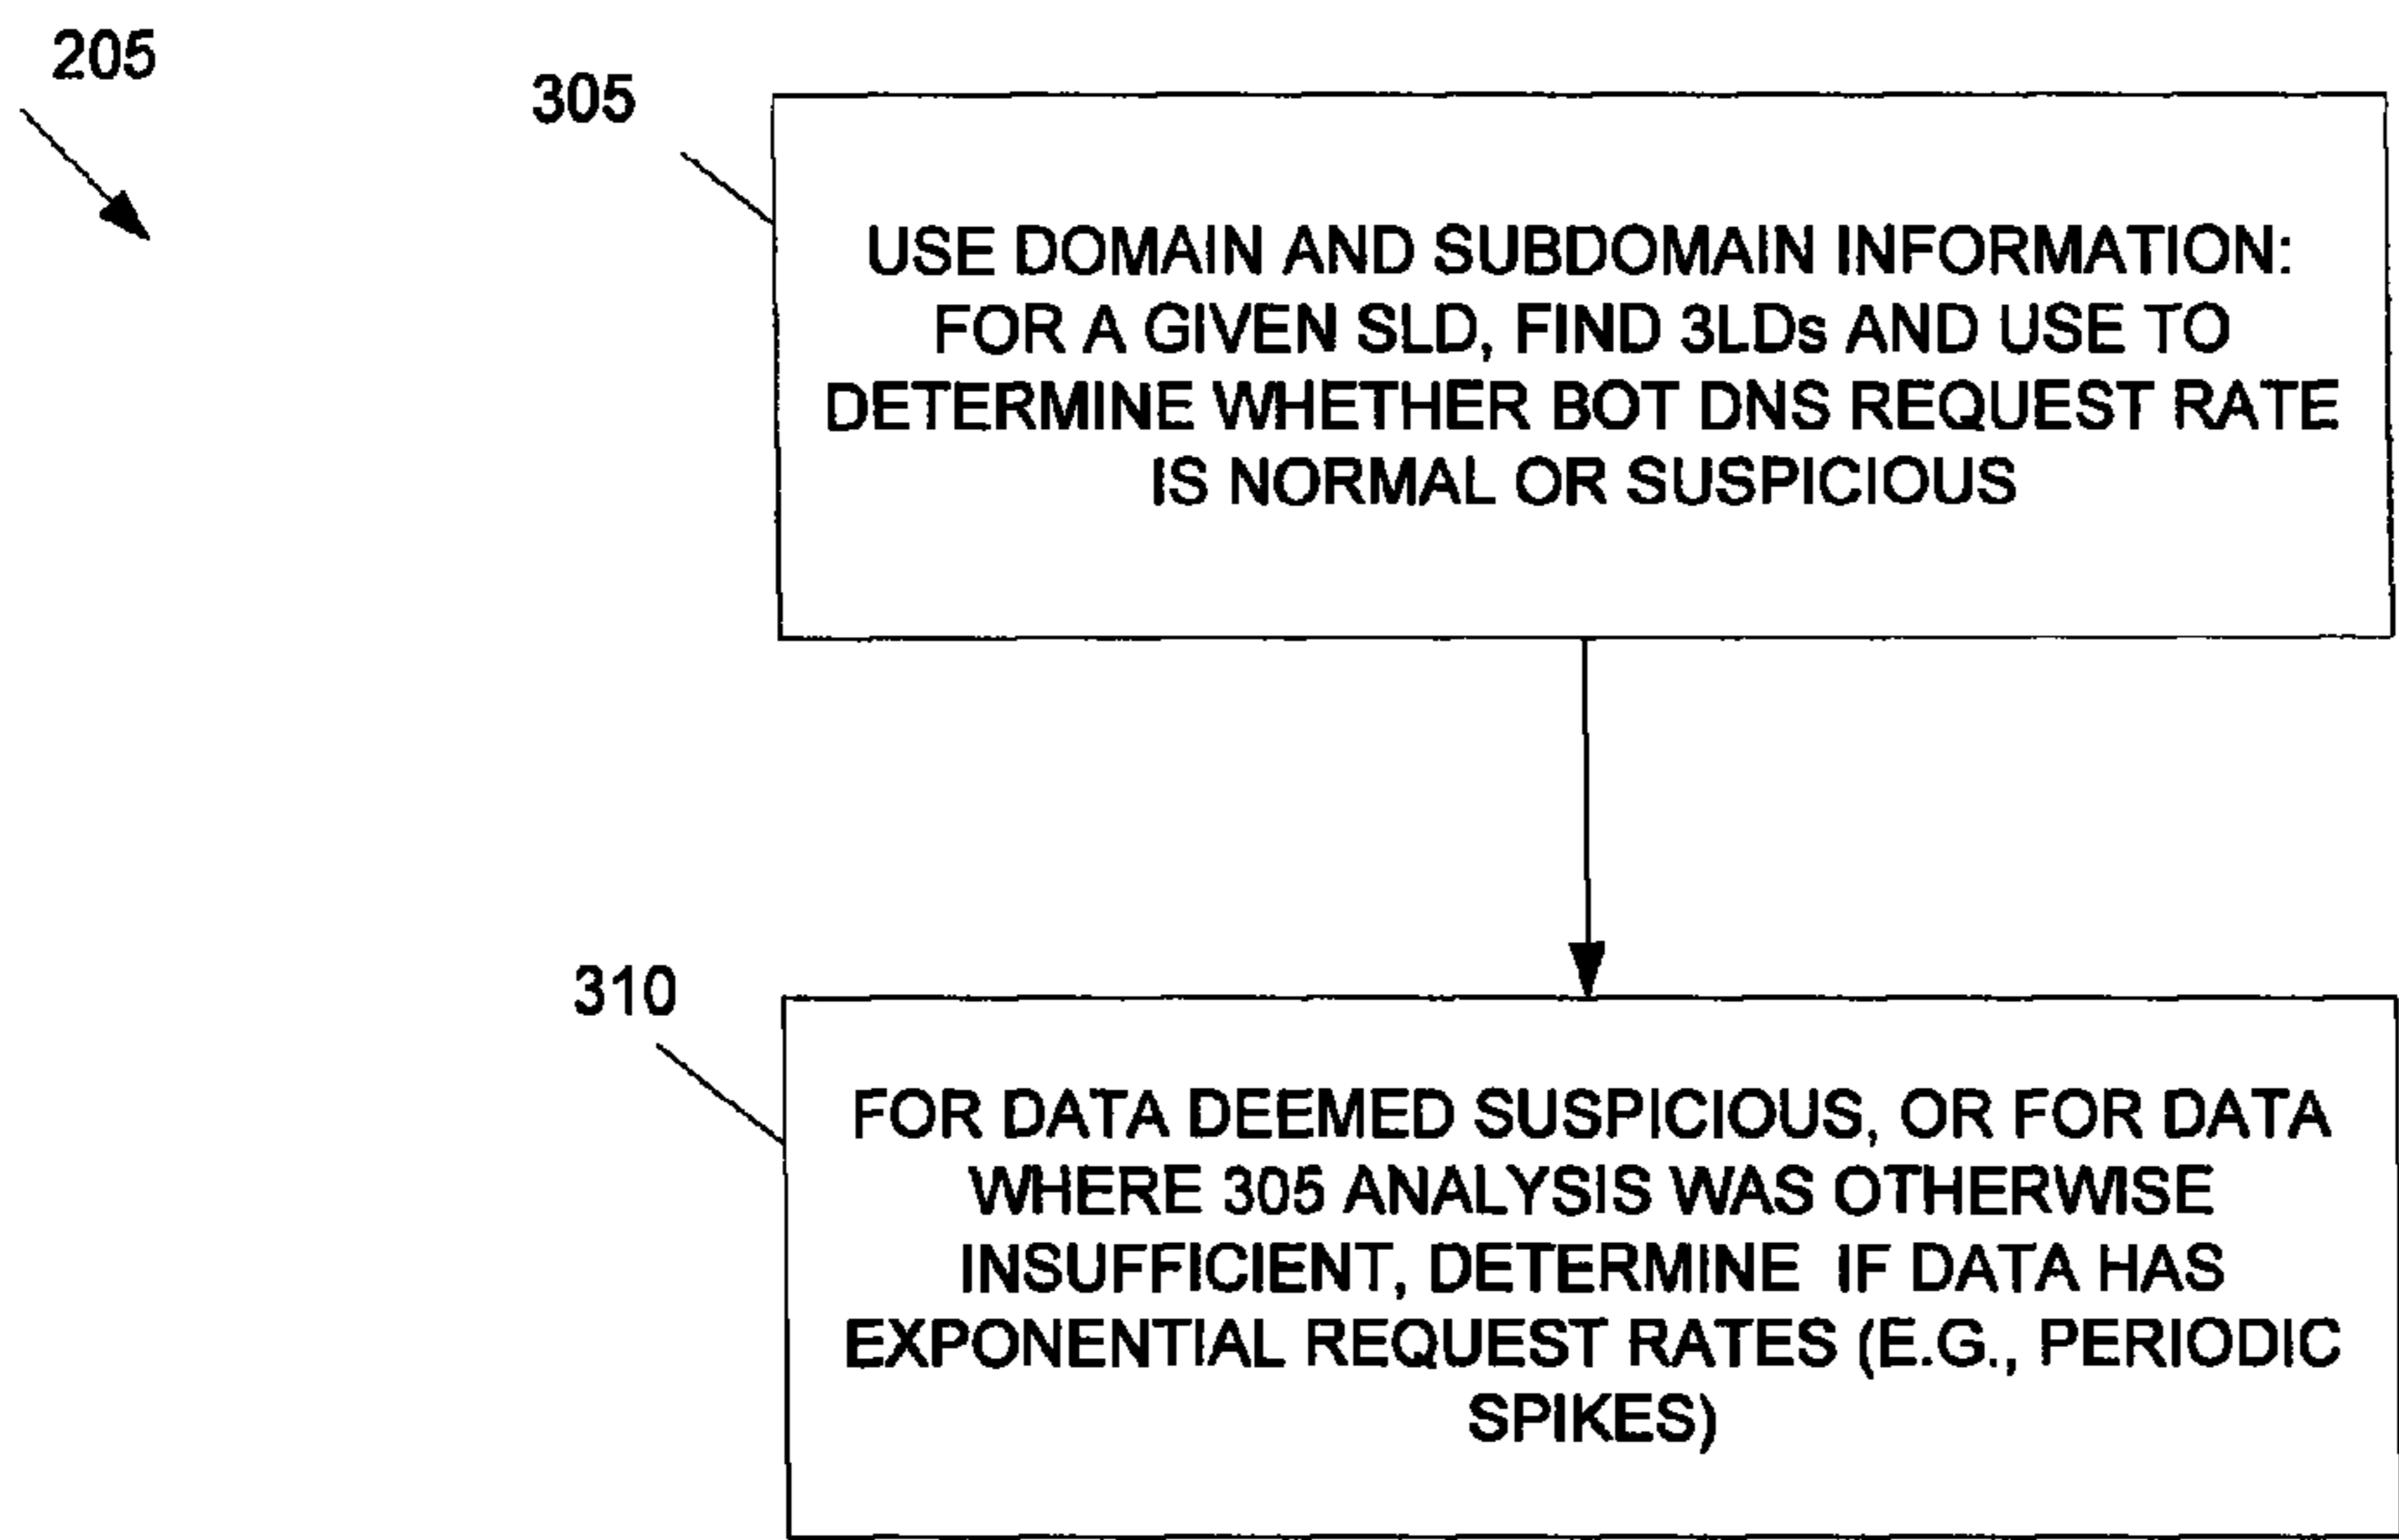


FIGURE 3



305 →

FIGURE 4A

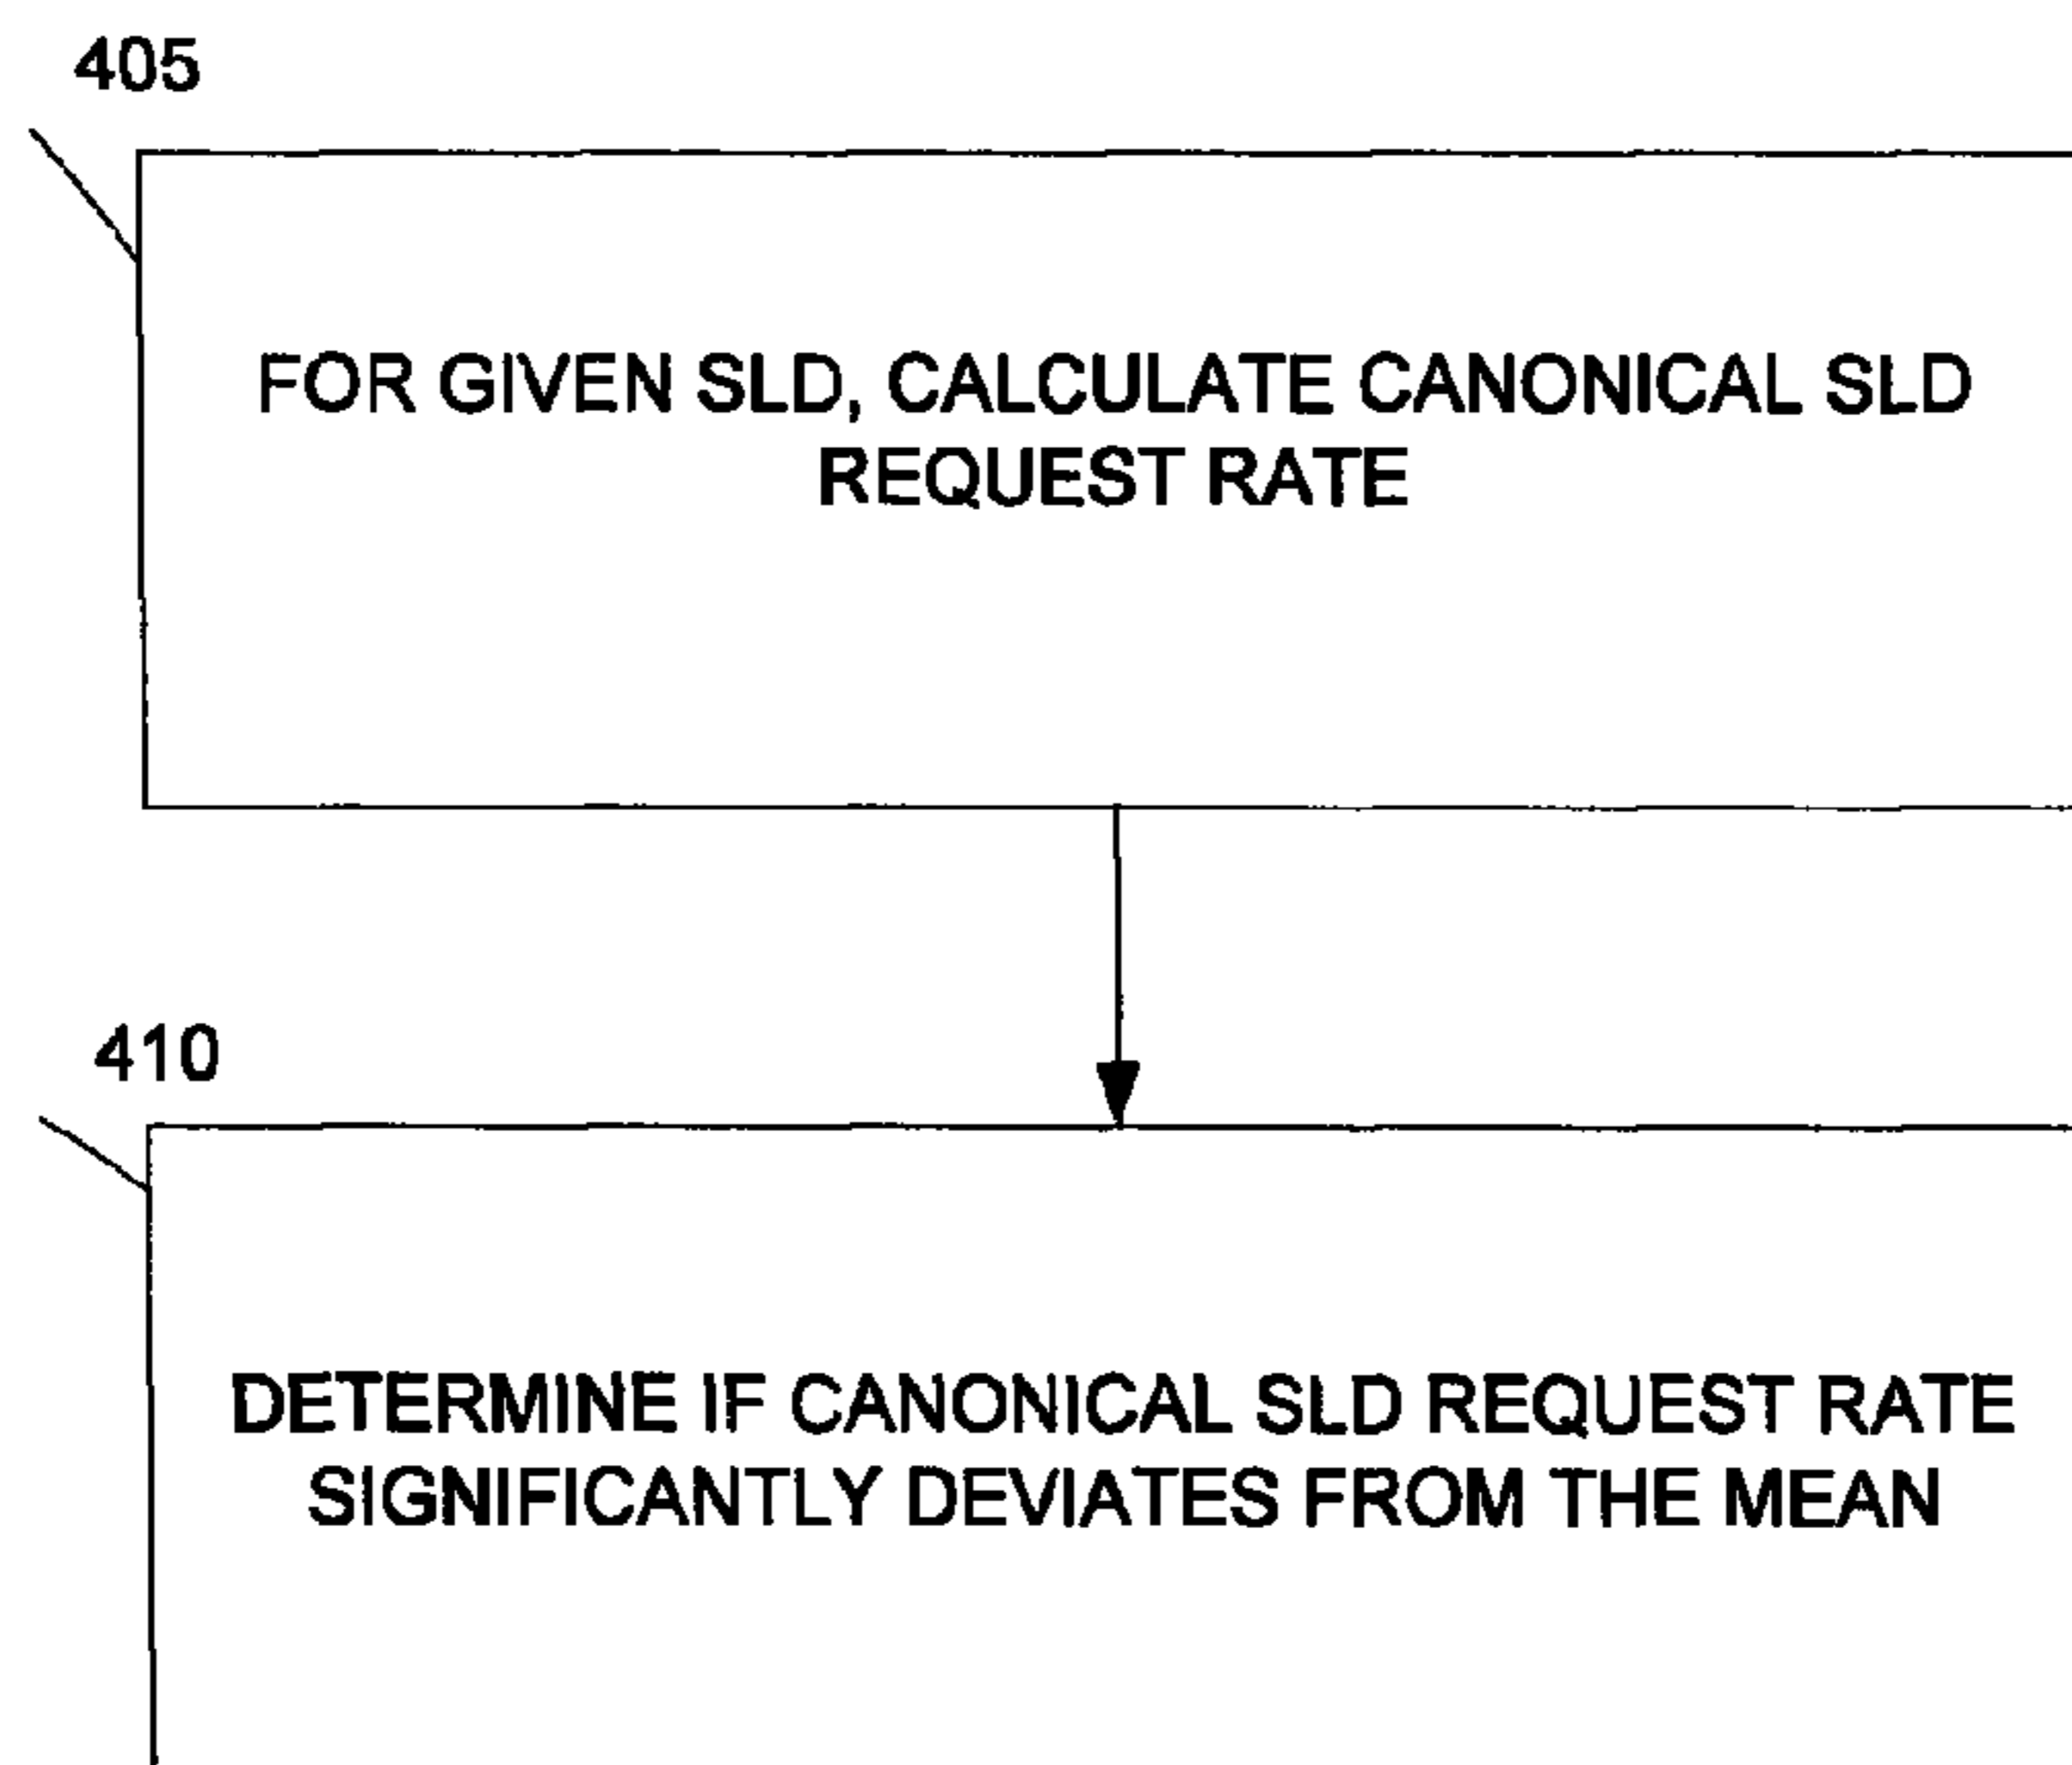


FIGURE 4B




FIGURE 5A




FIGURE 5B




FIGURE 6




FIGURE 7




FIGURE 8






FIGURE 9A




FIGURE 9B




FIGURE 10




FIGURE 11




FIGURE 12

<u>Node</u>	<u>ASN of Node</u>	<u>Out-Degree</u>	<u># of Children at Sinkhole</u>
1	Everyone's Internet (AS 13749)	36,875	12
2	Iquest (AS 7332)	32,159	7
3	UUNet (AS 701)	31,682	5
4	UPC Broadband (AS 6830)	26,502	8
5	E-xpedient (AS 17054)	19,530	4

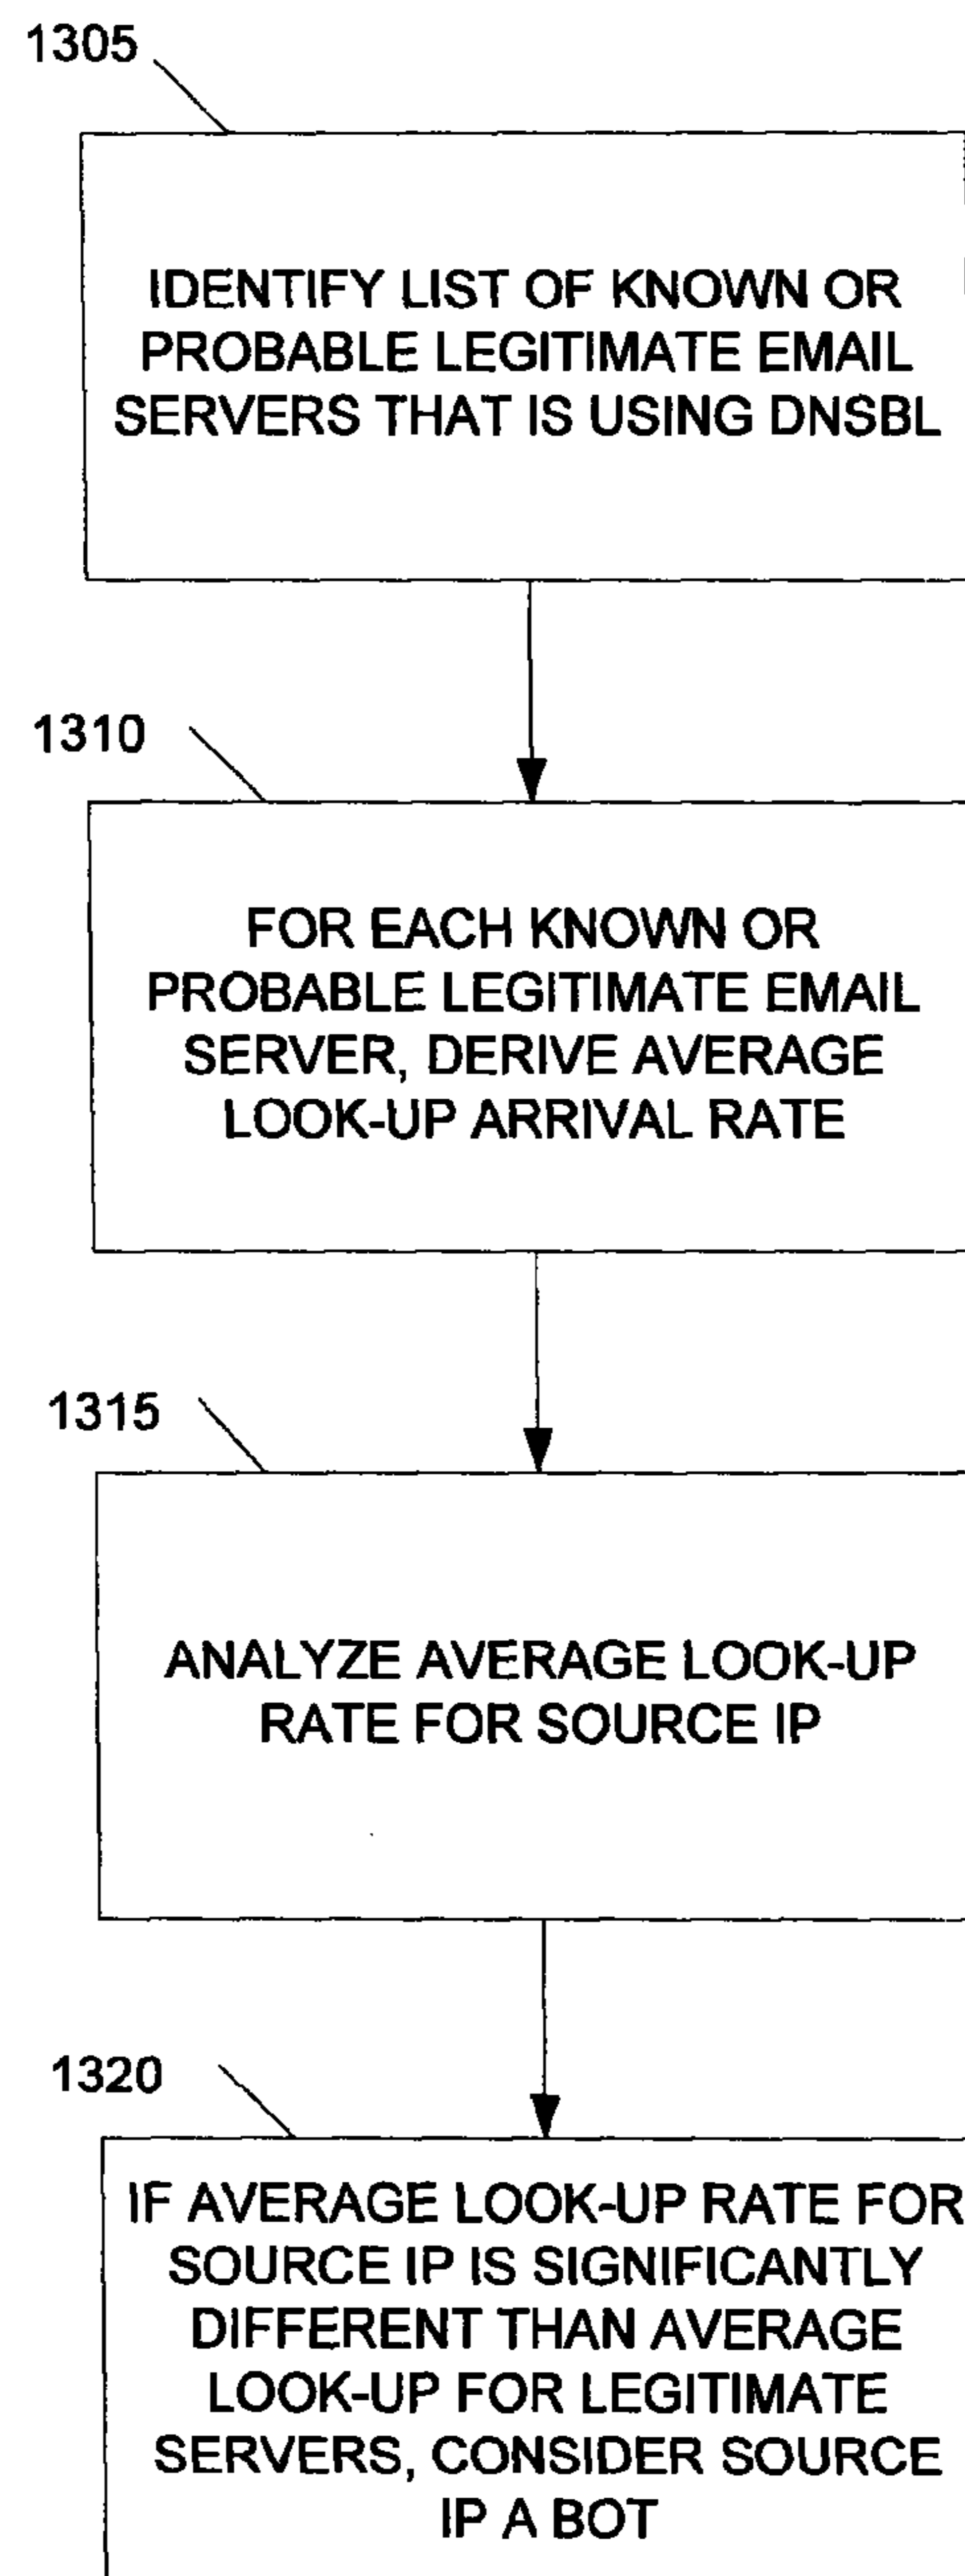
FIGURE 13

FIGURE 14

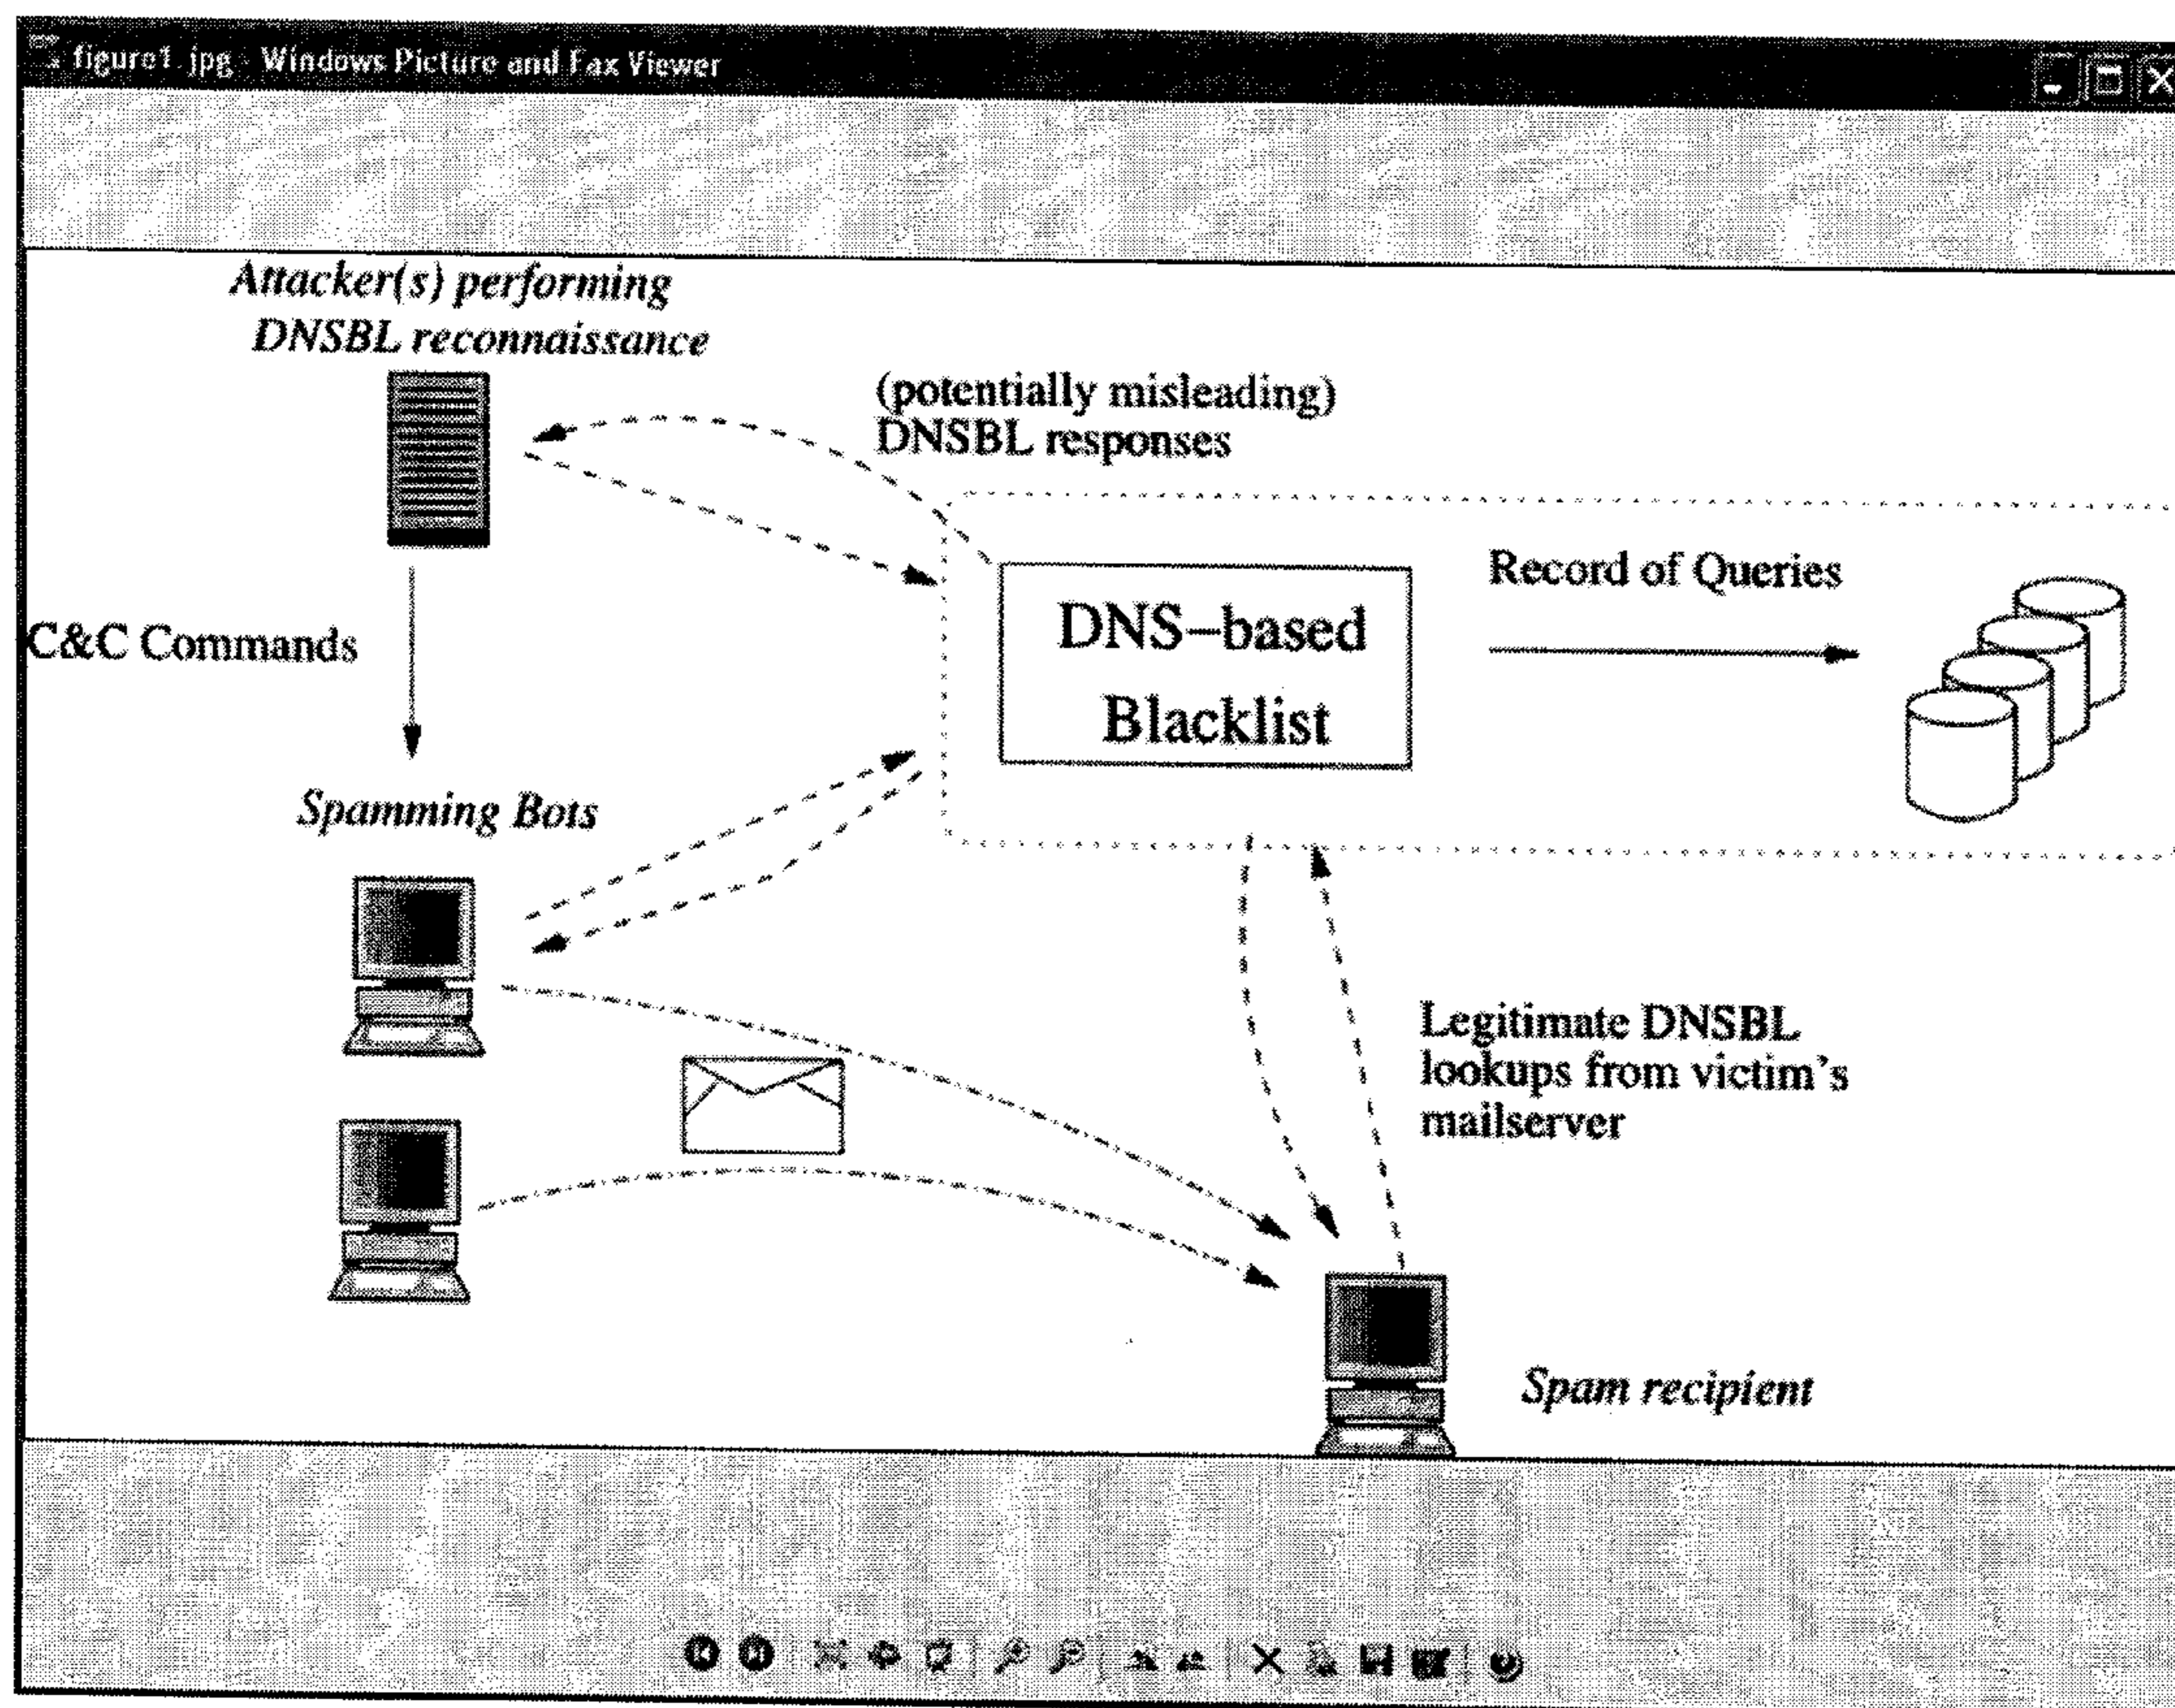


FIGURE 15




FIGURE 16

```
CONSTRUCTGRAPH()
create empty directed graph G

//Parsing
for each DNSBL query:
    Identify querier and queried

//Pruning
if querier in B or queried in B then
    add querier and queried to G if they
    are not already members of G
    if there exists an edge E (querier, queried) in G then
        increment the weight of that edge
    else
        add E (querier, queried) to G with weight 1
```

FIGURE 17




FIGURE 18




FIGURE 19

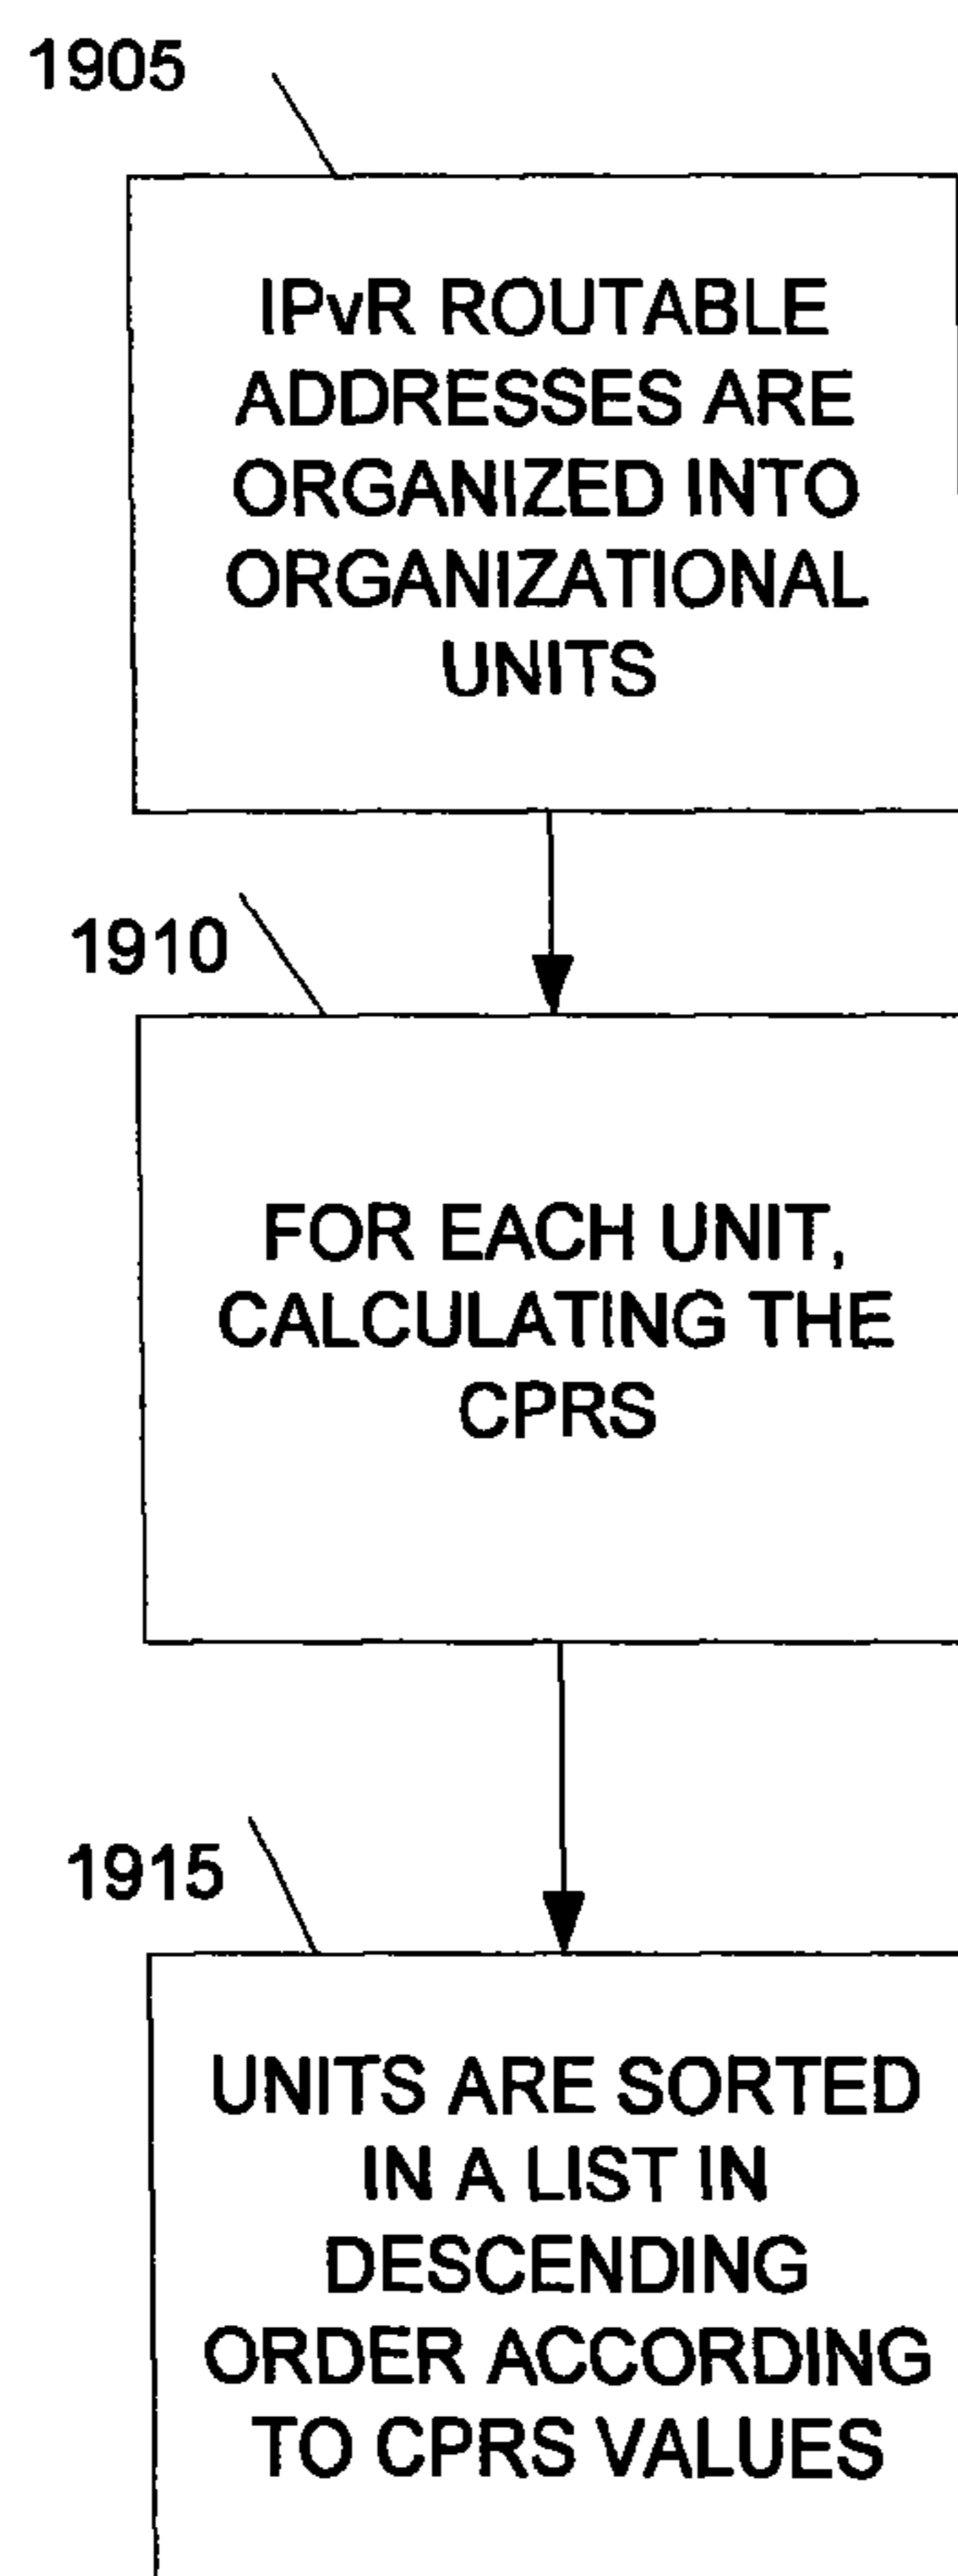


FIGURE 20




FIGURE 21

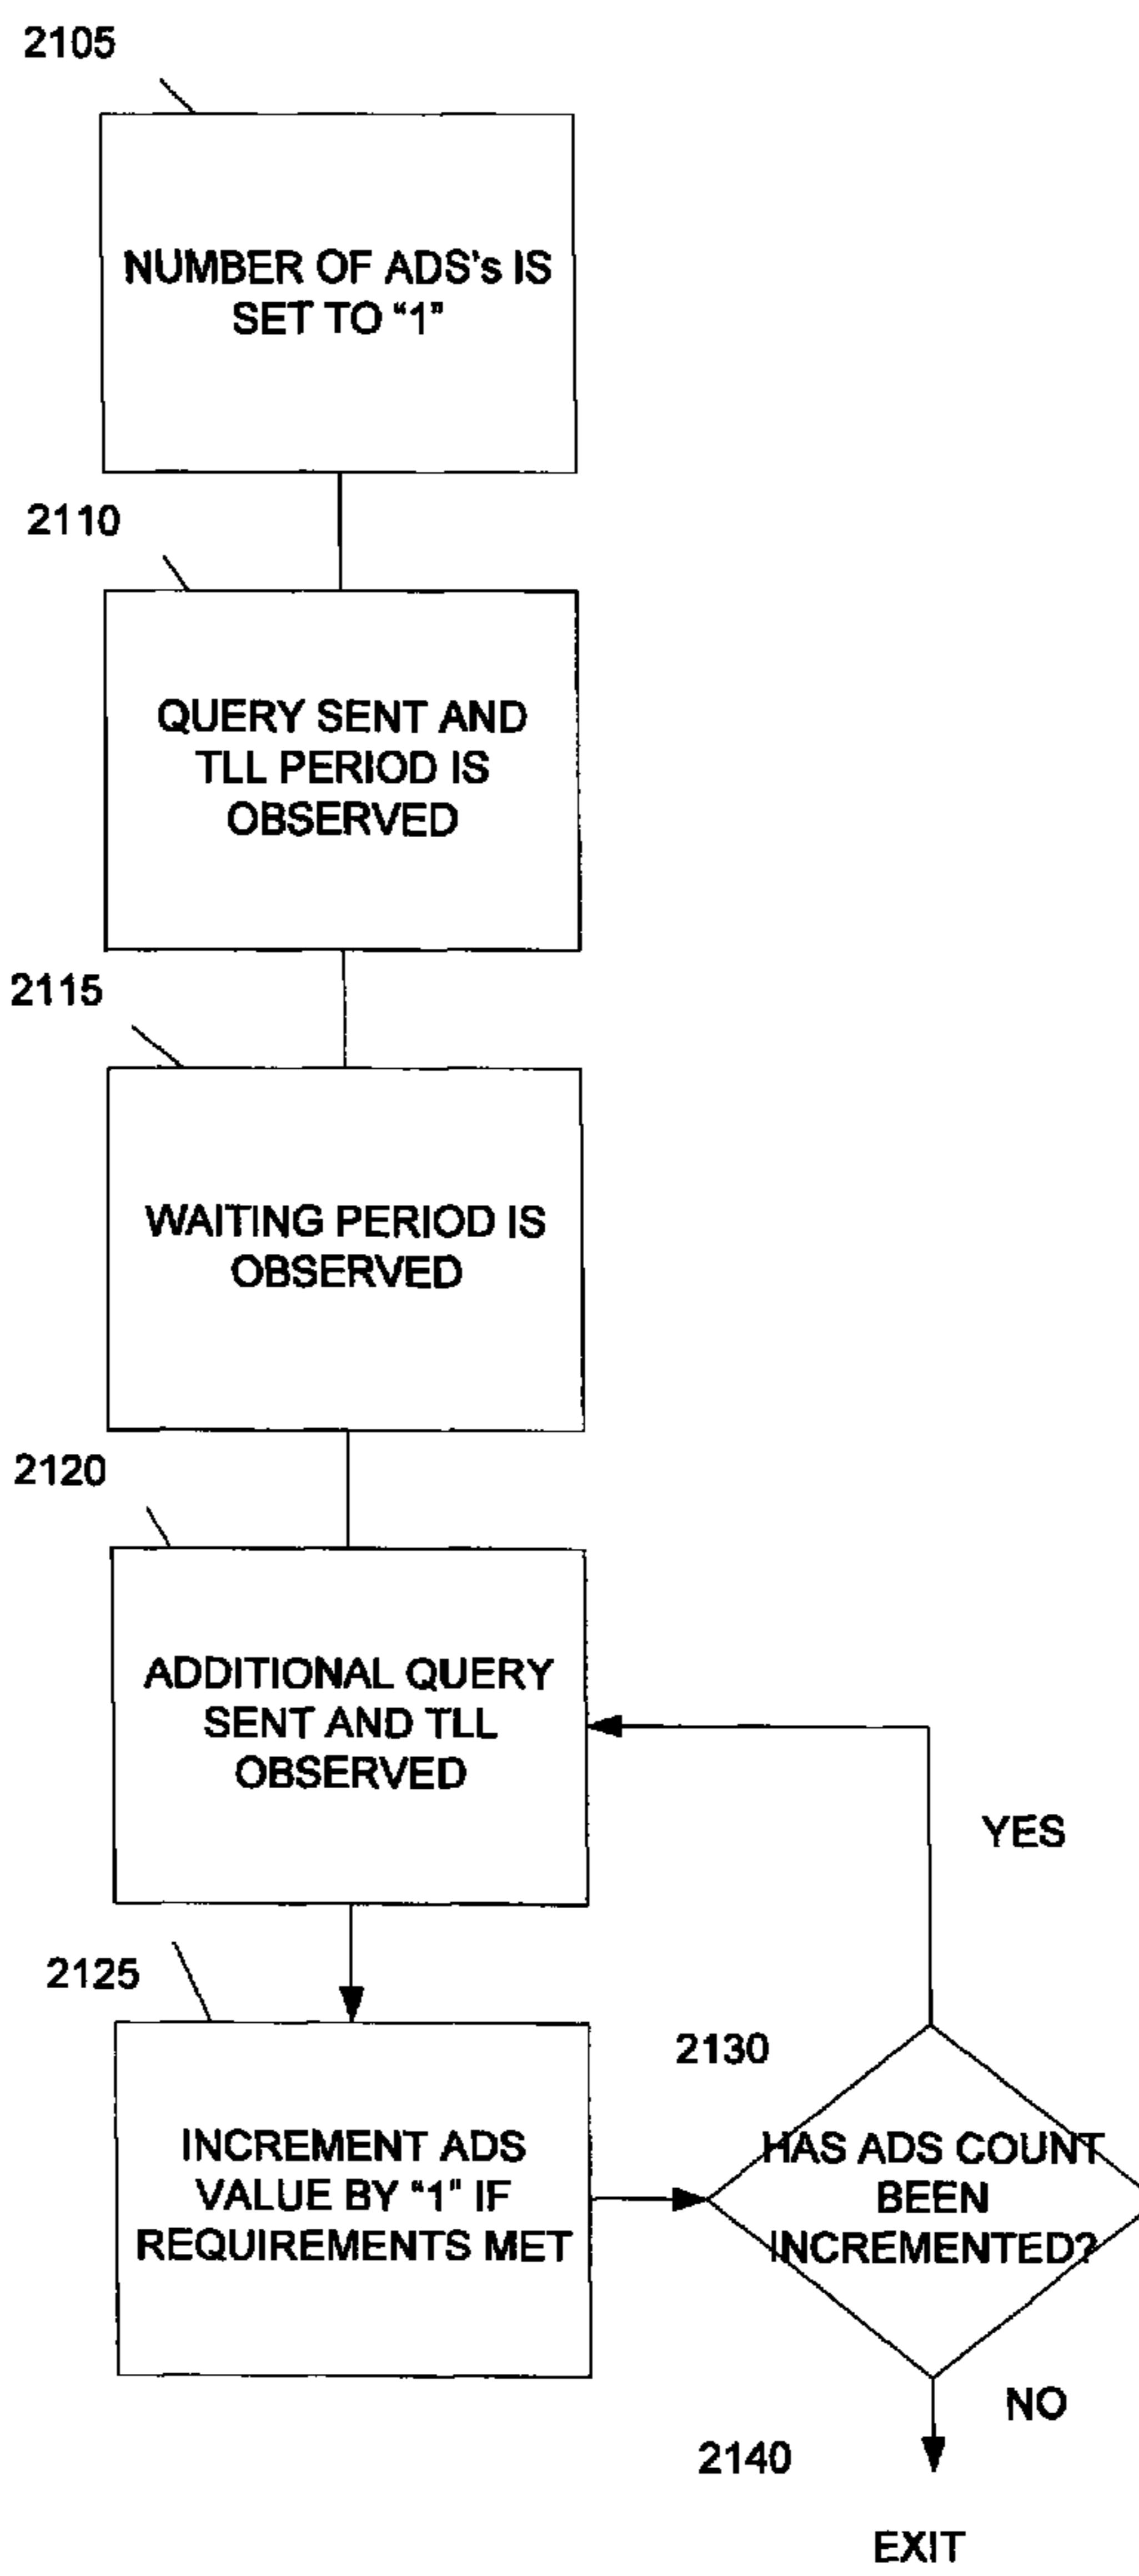
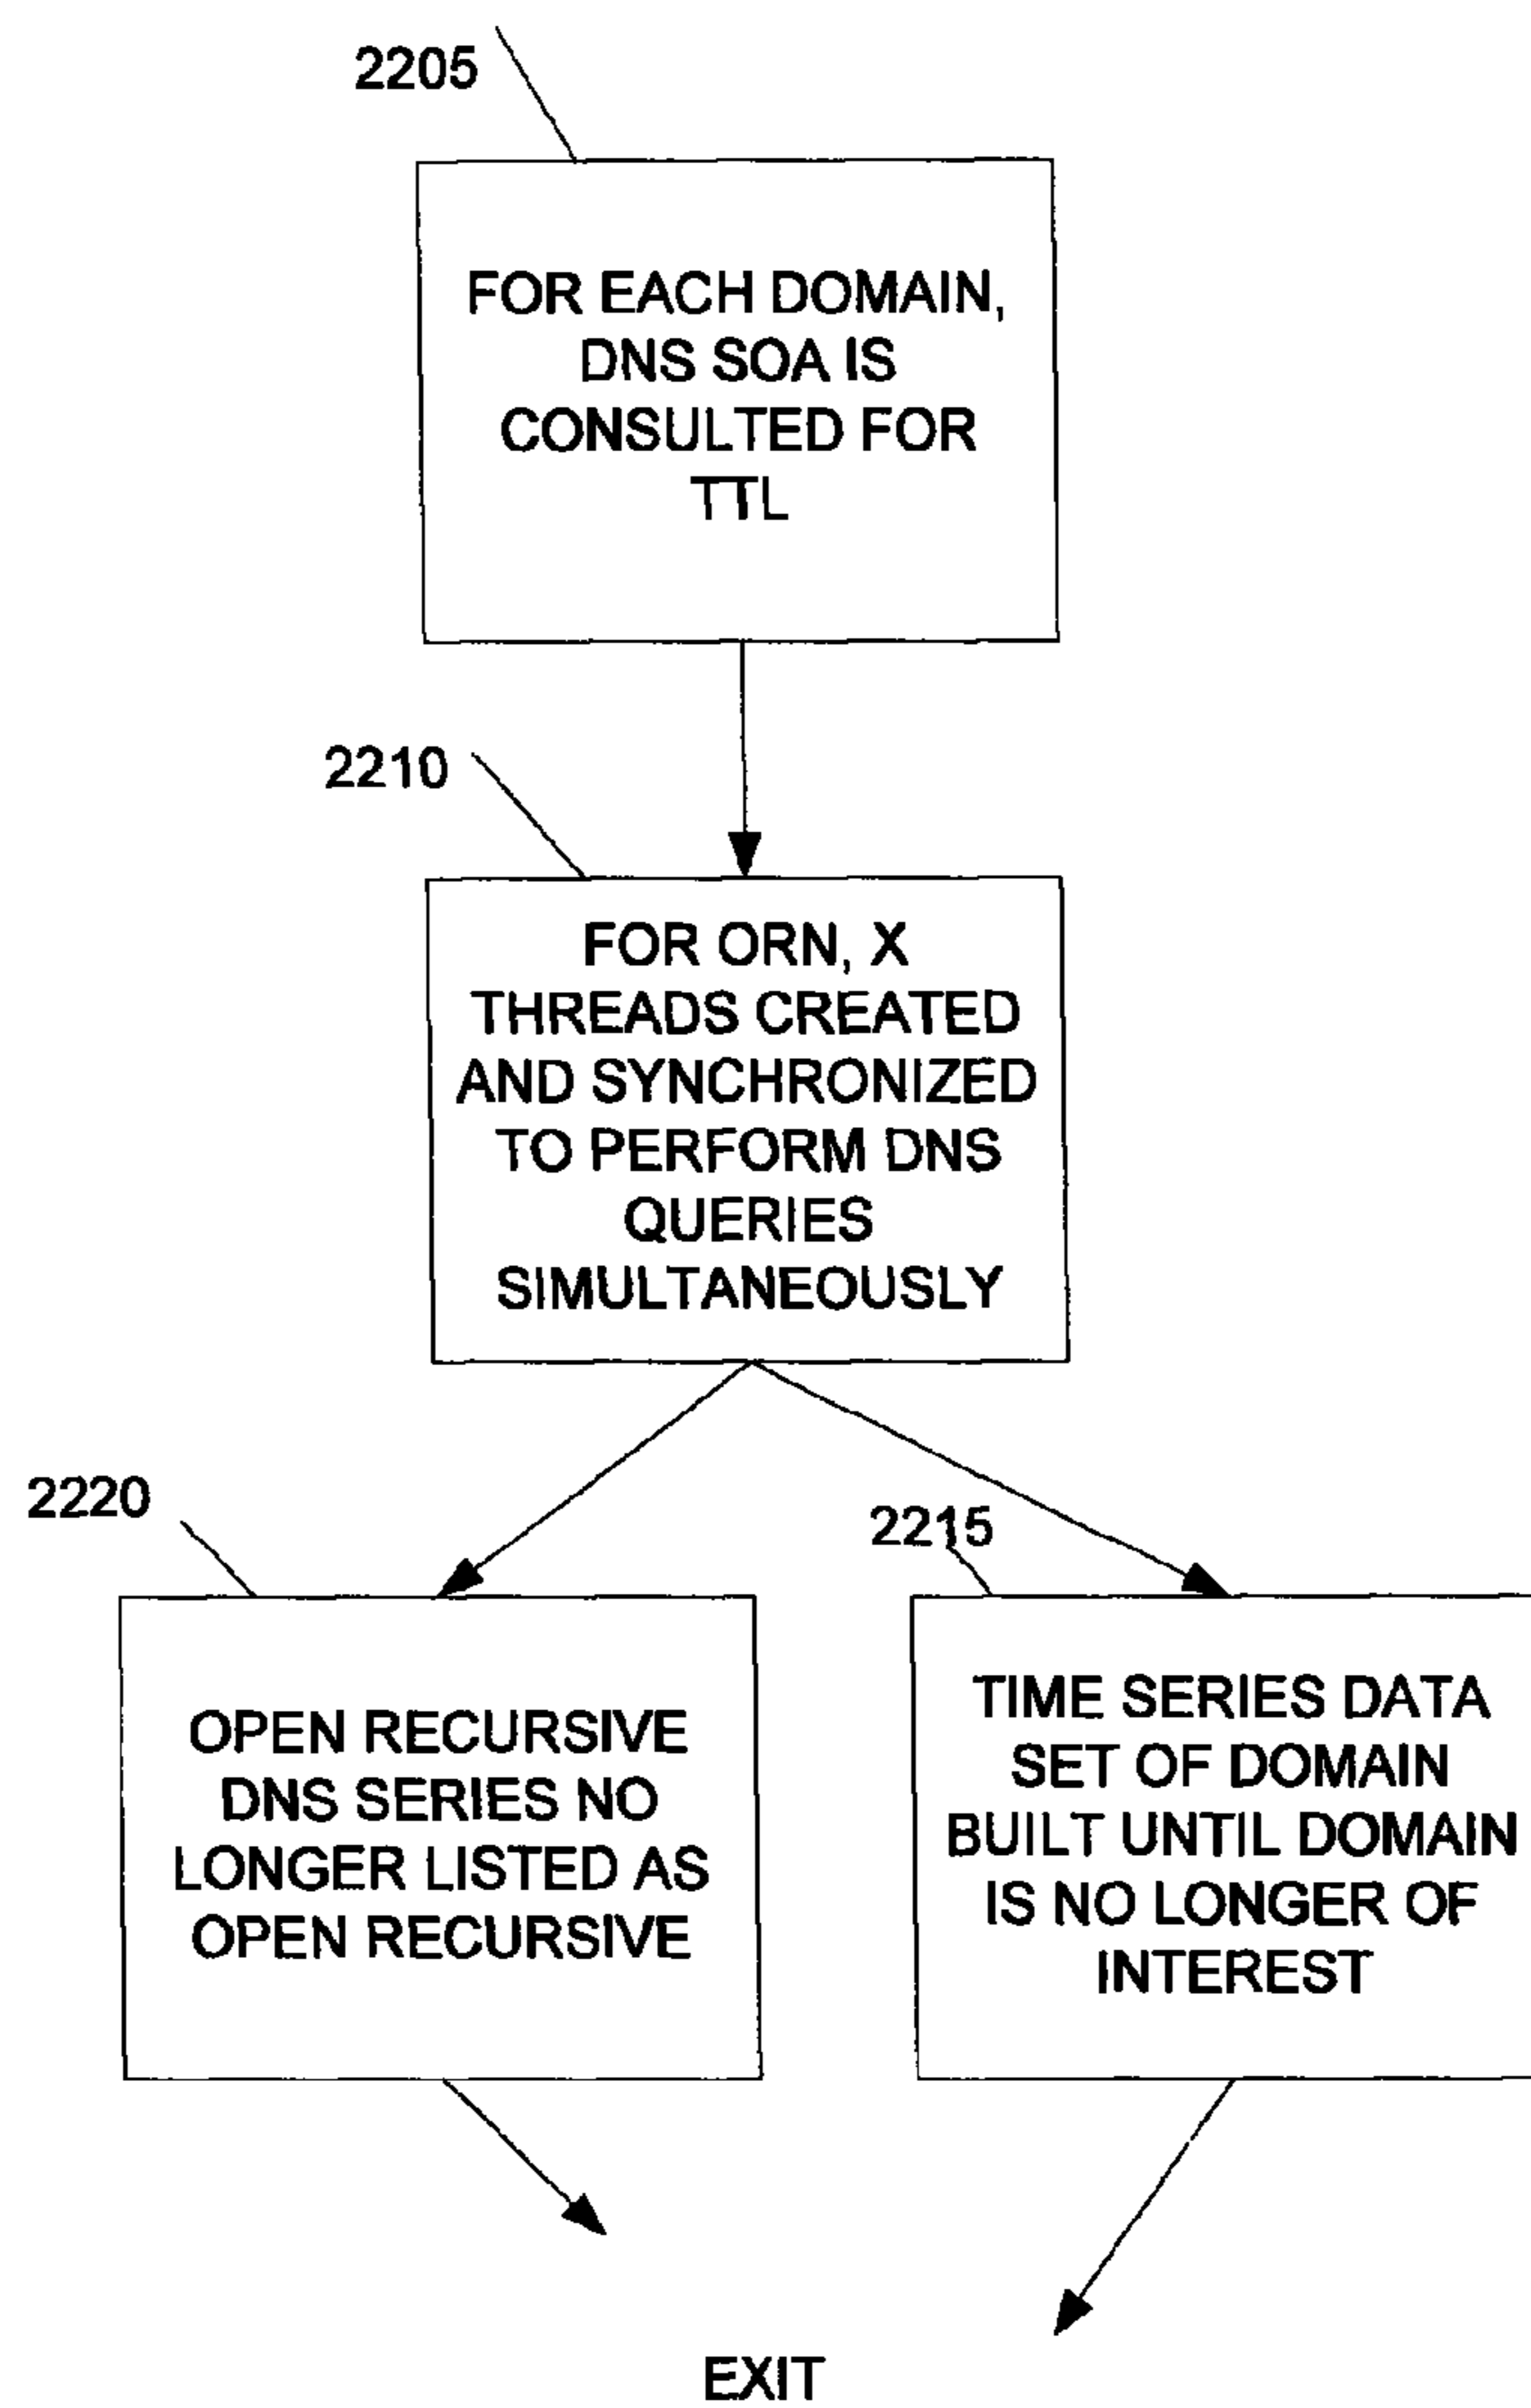


FIGURE 22



METHOD AND SYSTEM FOR DETECTING AND RESPONDING TO ATTACKING NETWORKS

This application claims priority to provisional application No. 60/730,615, entitled "Method to detect and respond to attacking networks," filed on Oct. 27, 2005, which is herein incorporated by reference. This application also claims priority to provisional application number 60/799,248, entitled "Revealing botnet membership using DNSBL counter-intelligence," filed on May 10, 2006, which is also herein incorporated by reference.

This application is supported in part by NSF grant CCR-0133629, Office of Naval Research grant N000140410735, and Army Research Office contract W911NF0610042.

BRIEF DESCRIPTION OF THE DRAWINGS

FIGS. 1A and 1B illustrates a system and method for botnet creation.

FIGS. 2A-9B illustrate several methods of detecting and disrupting botnets using DNS monitoring and sinkholing, according to several embodiments of the invention.

FIGS. 10-16 illustrate several methods for detecting and disrupting botnets using DNSBL monitoring, according to several embodiments of the invention.

FIGS. 17-22 illustrates methods for detecting and disrupting botnets using DNS cache snooping, according to several embodiments of the invention.

DESCRIPTION OF EMBODIMENTS OF THE INVENTION

Dynamic DNS Monitoring and Sinkholing

In one embodiment, the present invention is a method and system for identifying and/or attacking botnets. A bot is a robot or compromised computer that is used to carry out an attack. Examples of attacks include, but are not limited to, Distributed Denial of Service (DDOS) attacks, hosting distributed phishing pages, and key cracking. A botnet is a collection of bots. Botnets are composed of the bot victims reaped from different viruses, worms and Trojans. Thus, botnets are often referred to as viruses, worms or Trojans, depending on the context. The original infections compel the victim bots to run bot programs, which allow for remote administration.

Botnet Creation

To better understand how to detect and respond to botnets, an example pattern of botnet creation is presented in FIGS. 1A and 1B. FIG. 1A illustrates a system of botnets utilized in an attack. FIG. 1A illustrates a malware author **5**, a victim cloud of bot computers **10**, a Dynamic Domain Name System (DDNS) server **15**, and a Command & Control (C&C) computer **25**. Upon infection, each bot computer **10** contacts the C&C computer **25**. The malware author **5** (i.e., a hacker, denoted as VX) uses the C&C computer to observe the connections and communicate back to the victim bot computers **10**. Often, more than one C&C computer **25** is used. If not, a single abuse report can cause the C&C computer **25** to be quarantined or the account suspended. Thus, malware authors use networks of computers to control their victim bot computers **10**. Internet Relay Chat (IRC) networks are often utilized, as they are very resilient, and designed to resist hacker attacks. Because many public IRC networks are now patrolled by hacker-resistant software, botnets are migrating to private, non-IRC compliant services. In addition, malware authors **5** often try to keep their botnets mobile by using

DDNS service **15**, a resolution service that facilitates frequent updates and changes in computer locations. Each time the botnet C&C computer **25** is shut down by authorities, the botnet authors merely create a new C&C computer **25**, and update the DDNS entry. The bot computers **10** perform periodic DNS queries and migrate to the new C&C location. This practice is known as bot herding.

FIG. 1B illustrates a method of utilizing botnets for an attack. In **105**, the malware author **5** (e.g., VX) purchases one or more domain names (e.g., example.com), perhaps using a stolen account. The newly purchased domain names are initially parked at 0.0.0.0 (reserved for unknown addresses). A DNS or DDNS service can be used, in one embodiment. In **115**, the malware author **5** hard-codes the purchased domain names into dropper programs, which are sent to the victim bot computers **10** so that the victim bot computers **10** will contact the domain name servers. The dropper programs are programs that have been designed or modified to install a worm and/or virus onto a victim bot computer **10**. In **120**, the malware author **5** creates a C&C computer **25** for victim bot computers **10** to use to communicate. The C&C computer **25** can be, for example, a high-bandwidth compromised computer, or a high-capacity co-located box. The C&C computer **25** can be set up to run an IRC service to provide a medium for the bots to communicate. Note that other services can be used, such as, but not limited to: web services, on-line news group services, etc. In **125**, the malware author **5** will arrange for DNS resolution of domain name and register with DDNS service **15**. The IP address provided for in the registration is for the C&C computer **25**. As DNS propagates, more victim bot computers **10** join the network, and within a day, the hot army swells. The victims who contact the C&C computer **25** are compelled to perform a variety of tasks, such as, for example, but not limited to: updating their Trojans, attacking other computers, etc. When a DDNS server revokes a contract for DNS service, the malware author **5** (i.e., botmaster) just moves on, and secures DNS from yet another company. If the co-location service revokes the C&C contract (or cleans the box, in the case where the malware author **5** has used a compromised C&C computer **25**), the malware author **5** just rents or steals another C&C computer **25**.

Detecting Botnets

FIGS. 2A-2C illustrate a system and method of detecting and disrupting the communications between botnets and their victim bot computers **10**.

FIG. 2C illustrates A system for detecting a first network of compromised computers in a second network of computers, comprising: a computer including DNS detection software **265**, adapted to be connected to a network **250** and DNS data for the network **250**, the DNS detection software **265** capable of: collecting DNS data for the network **250**; examining the collected data relative to DNS data from known comprised and/or uncompromised computers **235** in the network **250**; and determining the identity of compromised computers in the network **250** based on the examination.

FIG. 2A, as does FIG. 1A, illustrates a malware author **5**, a victim cloud of bot computers **10**, a Dynamic Domain Name System (DDNS) server **15**, and a Command & Control (C&C) computer **25**. However, FIG. 2A also includes a sinkhole computer **20**. The IP address of the C&C computer **25** is replaced with the IP address of the sinkhole computer **20**. The sinkhole computer is used to hold traffic redirected from another computer. This way, the network of bot computers **10** is isolated from the C&C computer(s), and the botnet loses the ability to act as a coordinated group. Although it is also helpful to clean up the victim computers, this requires coor-

dination among different networks and can take time. However, disrupting the C&C can deal an immediate blow to the botnet.

FIG. 2B illustrates the method of detecting and disrupting the communications between botnets and their victim bot computers 10. In 205, the Command and Control (C&C) computer 25 of the botnet (network of attacking compromised computers) is identified, as explained below with respect to FIG. 3. In 210, the IP address of the C&C computer 25 is replaced with the IP address of the sinkhole computer 20. In 215, the bot computers 10 looking up the C&C computer 25 will be told to contact the sinkhole computer 20 instead. In 220, when a bot computer 10 contacts the sinkhole computer 20, the sinkhole computer 20 will record the IP address of the bot computer 10. In 225, traffic from the bot computers 10 to the sinkhole computer 20 can be utilized to detect and disrupt communications in the botnet.

FIG. 3 illustrates how a botnet's C&C computer can be identified. In 305, domain and subdomain information is used to determine whether a bot computer's DNS (Dynamic Name System) request rate is normal or suspicious. In 310, if the bot computer's DNS request rate is determined to be suspicious, it is determined if it has an exponential request rate (e.g., periodic spikes). In addition, the exponential request rate can also be utilized when the first filter of 305 is otherwise ineffective, such as, but not limited to, for analysis of low-and-slow spreading worms and/or viruses.

FIG. 4A illustrates the details of how the domain and subdomain information is used to determine whether a bot's DNS request rate is normal, as set forth above in 305. A DNS is a hierarchical system by which hosts on the Internet have both domain name addresses, such as "example.com", and IP addresses (such as 192.17.3.4). When a user types in a DNS name ("example.com"), a DNS application makes a DNS request by passing the DNS name and waiting for a response, such as the corresponding IP address or an error. DNS requests can be classified as either second-level domain (SLD) requests, such as "example.com", or third-level subdomain requests (3LD), such as "foo.example.com". To avoid increased costs and additional risks, botmasters often create botnets within 3LDs, all under a common SLD. For example, a botmaster may purchase the string "example.com" from a registrar, and then also purchase DDNS service for the 3LDs "botnet1.example.com", "botnet2.example.com", and so on. The botmasters use subdomains in order to avoid the purchase of a new domain name with each new botnet, e.g., "example1.com", "example2.com". Each purchase of a domain and name service involves risk. For example, the seller may be recording the originating IP for the transaction, and requiring the bot master to use numerous stepping stones. Some registrars are careful about screening and validating the "whois" contact information provided by the domain purchaser. If the purchase is performed with stolen user accounts, there is a further risk of being caught. Since many DDNS providers offer subdomain packages (e.g., a few free subdomains with DDNS service) this allows the botmaster to reuse their purchased domain and minimize both their costs and risk.

Botmasters also see another advantage in using subdomains. Even if service to a 3LD is suspended, service to other 3LDs within the same SLD is usually not disrupted. So, if "obtnet1.example.com" is sent to sinkhole computer, traffic to "normaluser.example.com" and "botnet2.example.com" is not disrupted. (Some DDNS providers may aggressively revoke accounts for the entire SLD, however, depending on

the mix of users.) This lets botmasters create multiple, redundant DDNS services for their networks, all using the same SLD.

By comparison, most normal users usually do not employ subdomains when adding subcategories to an existing site. For example, if a legitimate company owns "example.com" and wants to add subcategories of pages on their web site, they are more likely to expand the URL (e.g., "example.com/products") instead using a 3LD subdomain (e.g., "products.example.com"). This lets novice web developers create new content cheaply and quickly, without the need to perform complicated DNS updates (and implement virtual host checking in the web server) following each change to a web site.

Thus, normal users tend to have a single domain name (with subcategories of content hanging off the URL), while bot computers tend to use mostly subdomains. Of course, botmasters could decide to exclusively use SLDs for their botnets instead of 3LDs, but this doubles their cost (because each domain name must be purchased in addition to the original SLD) and increases the number of potentially risky financial transactions (that may lead to traceback) required to create the network.

Thus, to determine the number of 3LDs, in 405, for a given SLD, the canonical SLD DNS request rate is calculated. The canonical SLD request rate is defined as the total number of requests observed for all the 3LDs present in a SLD, plus any request to the SLD. We use the term $|SLD|$ to represent the number of 3LDs observed in a given SLD. Thus, if the SLD "example.com" has two subdomains "one.example.com" and "two.example.com", then its $|SLD|=2$. For a given SLD_i , with rate R_{SLD_i} , we calculate its canonical rate C_{SLD_i} as:

$$C_{SLD_i} = R_{SLD_i} + \sum_{j=1}^{|SLD_i|} R_{3LD_j}$$

where:

- R_{SLD_i} = SLD request rate
- R_{3LD_j} = 3LD request rate
- i = the SLD under consideration ($i=1, 2, \dots$)
- $j=1, 2, \dots$

Once the canonical SLD request rate is determined, in 410 it is determined if the canonical SLD request rate significantly deviates from the mean. When put in canonical form, distinguishing the normal and bot computer traffic is straight forward. The bottom line of FIG. 4B illustrates an average lookup rate for normal (i.e., non-bot) computers, in DNS requests per hour. An expected mean for the rate of normal traffic $E(X)=\mu$. Chebyshev's inequality is then used to fix an appropriate threshold for the normal request rates and request anomalies (i.e., bot) lookups. Chebyshev's inequality equation is:

$$P(|X - \mu| \geq t) \leq \frac{\sigma^2}{t}$$

where:

- P = the probability
- X = the rate of normal traffic
- μ = the mean of the rate of normal traffic
- t = the threshold
- σ = the standard deviation

The inequality places an upper bound on the chance that the difference between X and μ will exceed a certain threshold t.

5

As shown on the bottom line of FIG. 4B, normal traffic often uses only one SLD, and the traffic volume is low and relatively stable. In comparison, as shown on the upper line of FIG. 4b, botnets usually use one or more busy subdomains, which tend to have spikes in traffic.

FIG. 5A illustrates the details of how it is determined if a bot's DNS request rate has an exponential request rate, as set forth above in 310. In other words, the DNS density signature is determined. This test can be used as a second detection layer which can be used if the first filter is not effective. For example, the first filter could be evaded by botmasters if they adjust their use of SLDs or vary their DNS request rates, and thus blend in with normal traffic. In addition, noisy networks make the first filter ineffective because short-term normal and bot DNS rates may be very similar. An administrator may decide to revoke DDNS service for a host that has one or more "spikes" of traffic. To reduce the chance of false positives, a second filter can be used to examine just the hosts who have excessive canonical SLD scores.

A distinguishing feature for this second filter is that botnet DNS request rates are usually exponential over a 24 hour period. The diurnal nature of bot behavior means that there are periodic spikes in bot requests. These spikes are caused by infected hosts who turn on their computers in the morning, releasing a sudden burst of DNS traffic as the bots reconnect to the C&C computer. This spike is not present in normal DNS request rates, which require (usually slower and random) user interaction to generate a DNS request. In some cases, flash crowds of users visiting a popular site may behave like a botnet, but this is rare, and likely not sustained as seen in botnets.

Turning to FIG. 5A, in 505, the DNS request rates are sorted per hour. These sorted rates of normal DNS requests over a 24 hour period create a distribution, or density signature, for normal traffic. FIG. 5B illustrates sorted 24-hour average rates for normal traffic, as compared with sorted botnet traffic. The normal traffic is the bottom line, and the botnet traffic is the top line of FIG. 5B. Because of the diurnal spikes in traffic, the botnet traffic exhibits an exponential distribution.

Turning to 510, it is then determined if the sorted 24-hour traffic has any exponential activity. Any standard distance metric can compare the distributions. For example, the Mahalanbis distance can be used to measure the distance between request rate distributions and a normal model. (Note that other distance metrics can also be used.) The Mahalanobis distance, d , is:

$$d^2(x, \bar{y}) = (x - \bar{y})^T C^{-1} (x - \bar{y})$$

where:

x, \bar{y} = variable vectors (features) of the new observation and the trained (normal) profile

C = inverse covariance matrix for each member of the training data set

The Mahalanobis distance metric considers the variance of request rates in addition to the average request rate. This detects outliers, and measures the consistency of the observed request rates with the trained (normal) samples. The Mahalanobis distance metric can be simplified by assuming the independence of each sample in the normal traffic, and therefore removing the covariance matrix:

6

$$d(x, \bar{y}) = \sum_{i=0}^{n-1} \frac{|x_i - \bar{y}_i|}{\sigma_i}$$

where:

x, \bar{y} = variable vectors (features) of the new observation and the trained (normal) profile

n = the number of dimensions in the variable vectors

σ = the standard deviation

As with the canonical SLD request rate, training can be done using the normal model, and an appropriate threshold can be picked. Training can be done with a model of normal data, and a threshold chosen so that false positives are not generated. If observed traffic for a host has too great a distance score from the normal, it is deemed an outlier, and flagged as a bot computer.

Because of the underlying diurnal pattern driving bot computer name lookups, the sorted request rates only become distinct when grouped into clusters at least several hours in length. For this reason, this secondary detection system can also be used for low-and-slow spreading worms, and as an additional filtration step for noisy networks.

25 Disrupting Botnets

FIG. 6 illustrates several response options once a bot computer is detected, as set forth above in 225 of FIG. 2B: surveillance reporting 605, DDNS removal 610, and tarpits 615. Surveillance reporting 605 merely records the traffic. The sinkhole passively gathers information about attacking networks in a database, and keeps records on victims, activities, OS type/patch levels, and other relevant information. This data is shared with others, including individuals responsible for network response, Border Gateway Protocol (BGP) routing, and other network maintenance. Infection reports can be issued to monitored networks, or can be used to augment other intrusion detection systems, and assist law enforcement investigations. In addition, infection reports can be used to rank the number of infected individuals within Classless Interdomain Routing (CIDR) blocks and Internet Service Providers (ISPs) for a "bot reputation" score, refusing Simple Mail Transfer Protocol (SMTP) sessions from bot computers (to decrease spam), detecting click fraud in online advertising, or other research.

Another response option, DDNS removal 610, is to simply remove the botnets DDNS entry or name registration. Once the traffic is deemed abusive, and measured in the sinkhole, it is possible to revoke the DDNS account. Moreover, it is also possible in some cases to revoke the domain registration used by a botnet. Registration can be revoked where "whois" contact information is missing or proven wrong.

An additional optional response is the use of tarpits 615. There are at least two general types of tarpits: network layer (playing "TCP games") and application layer (honeypots). For network tarpits, in response to incoming bot synchronous (SYN) requests, bots can be sent a reset (RST), blackholed (i.e., given no response), sent a single acknowledgment, given multiple acknowledgments, or handed off to different types of tarpits. Routing layer (LaBrae-style) tarpits, for example, are easily evaded by modern multi-threaded bots. Many bot computers blacklist Internet Protocols (IPs) that repeatedly timeout or behave like a tarpit. Other bot computers use special application layer protocols or port-knocking (i.e., finding ports that are open) to detect tarpits and rival (hijacking) C&C computers.

For this reason, network-level tarpits are not completely effective against all classes of bot computers. For bot com-

puters that have learned how to evade network-layer tarpits, an application-level tarpit is utilized. Many of these bot computers leave the non-application level sinkhole because they expect a particular set of packets from the C&C computer, such as a port-knocking sequence or special banner message from an Internet Relay Chat (IRC) server. A limited proxy can be used to learn the appropriate hand-shake login sequence the bot expects. The bot computers first join the sinkhole, and are sent to an application-layer tarpit, also called a honeypot. The honeypot sends a "safe" heuristic subset of commands to the C&C computer, and observes the proper response behavior. Unsafe instructions (e.g., commands to scan networks or download other malware) are discarded, since this might expose a bot computer to instructions encoded in the channel topic. Even custom-made, non-RFC compliant protocols, such as heavily modified IRC servers, cannot evade application sinkholing, which slowly learns the proper sequence of instructions to fool the bot computers.

Analyzing Botnets

Modeling Prior Botnets to Predict Future Botnets.

In addition to the responses explained above, experience with previous botnets can also be used to predict the behavior of future botnets. Botnets are very widespread, so it is helpful to comparatively rank them and prioritize responses. Short-term variations in population growth can also be predicted, which is helpful because most dropper programs are short lived. In addition, different botnets use a heterogeneous mix of different infections exploiting different sets of vulnerabilities, often in distinct networks, with variable behavior across time zones. A model that can express differences in susceptible populations, and gauge how this affects propagation speed, is useful.

Botnets have a strongly diurnal nature. FIG. 7 illustrates a plot of SYN rates over time, broken down by geographic regions. A SYN rate is the rate of connection requests. The diurnal nature is likely because many users turn their computers off at night, creating a natural quarantine period, and varying the number of victim computers available in a geographical region. Such significant changes in populations over time affects propagation rates. Thus, there are different propagation rates, depending on time zone and time of day. Time zones not only express relative time, but also geography. If there are variable numbers of infected hosts in each region, then the natural quarantine effect created by a rolling diurnal low phase can have a significant impact on malware population and growth. Thus, a model is utilized to express the variable number of infected hosts, time zones, and regions of the Internet. This model allows estimation of short-term population projections for a given work, based on its regional focus, and the time of day. The model illustrates when bot computers spread the fastest, and allow comparison of short-term virulence of two different bot computers. This in turn can be used to improved surveillance and prioritize responses.

As illustrated in FIG. 7, the computers in each time zone are modeled as a group. The computers in each time zone have the same diurnal dynamics, regardless of whether they are infected or still vulnerable. The diurnal property of computers is determined by computer user's behavior, not by the infection status of computers.

As the number of infected computers in a region varies over time, $\alpha(t)$ is defined as the diurnal shaping function, or fraction of computers in a time zone that are still on-line at time t . Therefore, $\alpha(t)$ is a periodical function with a period of 24 hours. Usually, $\alpha(t)$ reaches its peak level at daytime (when users turn on their computers) and its lowest level at night (when users shut off their computers).

Diurnal Model for Single Time Zone.

First, a closed network within a single time zone is considered. Thus, all computers in the network have the same diurnal dynamics. It should be noted that the diurnal property of computers is determined by computer user behavior (e.g., turning on the computer at the beginning of the day). For the formula below, $I(t)$ is defined as the number of infected hosts at time t . $S(t)$ is the number of vulnerable hosts at time t . $N(t)$ is the number of hosts that are originally vulnerable to the worm under consideration. The population $N(t)$ is variable since such a model covers the case where vulnerable computers continuously go online as a worm spreads out. For example, this occurs when a worm propagates over multiple days. To consider the online/offline status of computers, the following definitions are used.

$I'(t)=\alpha(t)I(t)$ =number of infected online hosts at time t

$S'(t)=\alpha(t)S(t)$ =number of vulnerable hosts at time t

$N'(t)=\alpha(t)N(t)$ =number of online hosts among $N(t)$

To capture the situation where infected hosts are removed (e.g., due to computer crash, patching or disconnecting when infection is discovered), $R(t)$ is defined as the number of removed infected hosts at time t . Thus:

$$\frac{dR(t)}{dt} = \gamma I'(t)$$

where

γ =removal parameter, since only online infected computers can be removed (e.g., patched)

Thus, the worm propagation dynamics are:

$$\frac{dI(t)}{dt} = \beta I'(t)S'(t) - \frac{dR(t)}{dt}$$

where:

$$S(t)=N(t)-I(t)-R(t)$$

β =pair-wise rate of infection in epidemiology studies.

Note that for internet worm modeling

$$\beta = \frac{\eta}{\Omega}$$

where:

η =worm's scanning rate

Ω =size of IP space scanned by the worm

Thus, the worm propagation diurnal model is:

$$\frac{dI(t)}{dt} = \beta \alpha^2(t) I(t) [N(t) - I(t) - R(t)] - \gamma \alpha(t) I(t)$$

This diurnal model for a single time zone can be used to model the propagation of regional viruses and/or worms. For example, worms and/or viruses tend to focus on specific geographic regions because of the language used in the e-mail propagation system. Similarly, worms have hard-coded exploits particular to a language specific version of an Operating System (OS) (e.g., a worm that only successfully attacks Windows XP Home Edition Polish). For these regional worms and/or viruses, the infection outside of a single zone is negligible and the infection within the zone can be accurately

modeled by the above formula. It should also be noted that it is possible to not consider the diurnal effect. To so do, $\alpha(t)$ is set equal to 1.

Diurnal Model for Multiple Time Zones.

Worms and/or viruses are not limited to a geographic region. Victim bots are usually spread over diverse parts of the world, but can be concentrated in particular regions, depending on how the underlying infections propagate. For example, some attacks target a particular language edition of an operating system, or use a regional language as part of a social engineering ploy. For example, there are worms and/or viruses that contain enormous look-up tables of buffer-overflows offset for each language edition of Windows. Similarly, many email spreading worms and/or viruses use a basic, pigeon English, perhaps to maximize the number of Internet users who will read the message and potentially open up the virus. These regional variations in infected populations play an important role in malware spread dynamics. Thus, in some situations it is useful to model the worm and/or virus propagation in the entire Internet across different time zones. Since computers in one time zone could exhibit different diurnal dynamics from the ones in another time zone, computers in each zone are treated as a group. The Internet can then be modeled as 24 interactive computer groups for 24 time zones. Since many of the time zones have negligible numbers of computers (such as time zones spanning parts of the Pacific Ocean), worm propagation can be considered in K time zones where K is smaller than 24. For a worm and/or virus propagation across different time zones, the worm propagation for time zone i is:

$$\frac{dI_i(t)}{dt} = \sum_{j=1}^K \beta_{ji} I_j(t) S_i'(t) - \frac{dR_i(t)}{dt}$$

which yields:

$$\frac{dI_i(t)}{dt} = \alpha_i(t) [N_i(t) - I_i(t) - R_i(t)] \sum_{j=1}^K \beta_{ji} \alpha_j(t) I_j(t) - \gamma_i \alpha_i(t) I_i(t)$$

where:

$N_i(t)$ = the number of online hosts at time t in time zone i ($i=1, 2, \dots, K$)

$S_i(t)$ = the number of vulnerable hosts at time t in time zone i

$I_i(t)$ = the number of infected online hosts at time t in time zone i

$R_i(t)$ = the number of removed infected hosts at time t in time zone i

Similarly, $N_j(t)$, $S_j(t)$, $I_j(t)$, $R_j(t)$ = the number of hosts in time zone $j=1, 2, \dots, K$

$\alpha_i(t)$ = diurnal shaping function for the time zone i

β_{ji} = pairwise rate of infection from time zone j to i

γ_i = removal rate of time zone i

For a uniform-scan worm and/or virus, since it evenly spreads out its scanning traffic to the IP space:

$$\beta_{ji} = \frac{\eta}{\Omega}, \forall i, j \in K$$

where:

n = the number of scans sent to the group from an infected host in each time unit;

Ω = the size of the IP space in the group

For worms that do not uniformly scan the IP space:

$$\beta_{ji} = \frac{\eta_{ji}}{\Omega_i}$$

where:

n_{ji} = the number of scans sent to group i from an infected host in group j in each time unit;

Ω_i = size of IP space in group i

Thus, when a new worm and/or virus is discovered, the above equation can be used by inferring the parameter β_{ji} based on a monitored honeypot behavior of scanning traffic. (Note that a honeypot is a computer set up to attract malicious traffic so that it can analyze the malicious traffic.) As noted above with reference to FIG. 6, many honeypot systems can observe all outgoing scans sent out by a trapped worm and/or virus. The worm's scanning target address distribution can therefore be inferred based on reports from multiple honeypots. Then η_{ji} can be derived based on the worm's scanning distribution.

Thus, as illustrated in FIG. 8, equations and graphs can be produced showing the different effect of a virus and/or worm in different time zones. FIG. 8 illustrates the number of SYN connections sent to the sinkhole per minute from each of a North American group, a Europe group, and an Asia group. Note that all the groups shown in FIG. 8 have diurnal (i.e., exponential) behavior at different times of the day. Note that the North American and Asian groups have more noise, likely because countries in these groups tend to span numerous time zones with large numbers of infected individuals, and China has one time zone for the entire country. In comparison, the European countries tend to occupy a single zone, and most victims are located in the western-most time zones.

The diurnal models tell us when releasing a worm will cause the most severe infection to a region or the entire Internet. For worms that focus on particular regions, the model also allows prediction of future propagation, based on time of release. A table of derived shaping functions can be built, which are based on observed botnet data and other heuristics (e.g., the exploit used, the OS/patch level it affects, country of origin). When a new worm and/or virus is discovered, the table for prior deviations can be consulted to forecast the short-term population growth of the bot, relative to its favored zone and time of release.

In addition, knowing the optimal release time for a worm will help improve surveillance and response. To identify an optimal release time, the scenario is studied where the worm uniformly scans the Internet and all diurnal groups have the same number of vulnerable population, i.e., $N1=N2=N3$. To study whether the worm's infection rate β affects the optimal release time, the worm's scan rate η (remember

$$\beta = \frac{\eta}{\Omega})$$

is changed. The study of optimal release times is useful because we can better determine the defense priority for two viruses or worms released in sequence. Viruses often have generational releases, e.g., worm.A and worm.B, where the malware author improves the virus or adds features in each

new release. The diurnal model allows consideration of the significance of code changes that affect $S(t)$ (the susceptible population). For example, if worm.A locally affects Asia, and worm.B then adds a new feature that also affects European users, there clearly is an increase in its overall $S(t)$, and worm.B might become a higher priority. But when worm.B comes out, relative when worm.A started, plays an important role. For example, if the European users are in a diurnal low phase, then the new features in worm.B are not a near-term threat. In such a case, worm.A could still pose the greater threat, since it has already spread for several hours. On the other hand, if worm.B is released at a time when the European countries are in an upward diurnal phase, then worm.B could potentially overtake worm.A with the addition of the new victims.

The diurnal models in FIGS. 9A and 9B exposes such a counter-intuitive result. FIG. 9A illustrates worm and/or virus propagation at different release times. In addition, FIG. 9B shows the number of infected hosts at various release times. Thus, as illustrated above, researchers and/or computer managers are able to calculate optimal release items for worms and therefore rank them based on predicted short-term growth rates. Examples of utilizing diurnal models include, but are not limited to: priority ranking (short and long term), patch management, and/or filtration management. In priority ranking, diurnal models help computer managers figure out which botnet needs to be addressed first because they are able to estimate the maximum number of infected individuals from each bot during each time of day. In patch management allows, diurnal models help a computer manager to prioritize patches. For example, if a computer manager knows that a virus related to Microsoft 2000 is impacting a certain number of users at a certain time, he can use this knowledge to prioritize patches performed related to other botnet threats. In filtration management, diurnal models help a computer manager to determine if certain connections should be refused during certain times. For example, if a computer manager knows that during a certain time, email traffic from China will be highly infected, he can use a filter or firewall to refuse that traffic during a certain time period.

DNSBL Monitoring

Another method of passively detecting and identifying botnets (i.e., without disrupting the operation of the botnet) is through revealing botnet membership using Domain Name System-based Blackhole List (DNSBL) counter-intelligence. DNSBL can be used to passively monitor networks, often in real-time, which is useful for early detection and mitigation. Such passive monitoring is discreet because it does not require direct communication with the botnet. A bot that sends spam messages is usually detected by an anti-spam system(s) and reported/recorded in a DNSBL, which is used to track IP addresses that originate spam. An anti-spam system gives a higher spam score to a message if the sending IP address can be looked up on a DNSBL. It is useful to distinguish DNSBL traffic, such as DNSBL queries, that is likely being perpetrated by botmasters from DNSBL queries performed by legitimate mail servers.

Bots sometimes perform look-ups (i.e., reconnaissance to determine whether bots have been blacklisted) on the DNSBL. For example, before a new botnet is put in use for spam, the botmaster of the new botnet or another botnet may look up the members of the new botnet on the DNSBL. If the members are not listed, then the new botnet, or at least certain bots, are considered “fresh” and much more valuable.

If the bot performing reconnaissance is a known bot, e.g., it is already listed on the DNSBL or it is recorded in some other botnet database (e.g., a private botnet database), then the new

botnet can be identified using the IPs being queried by the bot. Analysis can be performed at the DNSBL server, and for each query to the DNSBL, the source IP issuing the query can be examined, and the subject IP being queried can also be examined. If the source IP is a known bot, then the subject IP is also considered to be a bot. All of the subject IPs that are queried by the same source IP in a short span of time are considered to be in the same botnet.

If an unknown bot is performing reconnaissance, it must first be identified as a bot, and then the IPs it queries can also be identified as bots. DNSBL reconnaissance query traffic for botnets is different than legitimate DNSBL reconnaissance query traffic. FIG. 10 illustrates several methods for analyzing reconnaissance traffic, according to several embodiments of the invention.

Self-Reconnaissance

In 1005, self-reconnaissance is detected. To perform “self-reconnaissance”, the botmaster distributes the workload of DNSBL look-ups across the botnet itself such that each bot is looking up itself. Detecting such botnet is straightforward because a legitimate mail server will not issue a DNSBL look-up for itself.

Single Host Third-Party Reconnaissance

In 1010, single host third-party reconnaissance is detected. To explain third-party reconnaissance, a look-up model is provided in FIG. 11. FIG. 11 illustrates IP address A looking up IP address B, according to one embodiment of the invention. A line from node A to node B indicates that node A has issued a query to a DNSBL to determine whether node B is listed.

A legitimate mail server both receives and sends email messages, and hence, will both perform look-ups (for the email messages it received in) and be the subject of look-ups by other mail servers (for the email messages it sent out). In contrast, hosts performing reconnaissance-based look-ups will only perform queries; they generally will not be queried by other hosts. Legitimate mail servers are likely to be queried by other mail servers that are receiving mail from that server. On the other hand, a host that is not itself being looked up by any other mail server is, in all likelihood, not a mail server but a bot. This observation can be used to identify hosts that are likely performing reconnaissance: lookups from hosts that have a low in-degree (the number of look-ups on the bot itself for the email messages it sent out), but have a high out-degree (the number of look-ups the bot performs on other hosts) are more likely to be unrelated to the delivery of legitimate mail.

In single host third-party reconnaissance, a bot performs reconnaissance DNSBL look-ups for a list of spamming bots. The in-degree (d_{in}) should be small because the bot is not a legitimate mail server and it has not yet sent a lot of spam messages (otherwise it will have been a known bot listed in DNSBL already). Thus, a look-up ratio α_A is defined as:

$$\alpha_A = \frac{d_{out}}{d_{in}}$$

where:

α_A = the look-up ratio for each node A

d_{in} = the in-degree for node A (the number of distinct IPs that issue a look-up for A).

d_{out} = the out-degree for node A (the number of distinct IPs that A queries)

Thus, utilizing the above formula, a bot can be identified because it will have a much larger value of α than the legitimate mail servers. Single-host reconnaissance can provide

13

useful information. For example, once a single host performing such look-ups has been identified, the operator of the DNSBL can monitor the lookups issued by that host over time to track the identity of hosts that are likely bots. If the identity of this querying host is relatively static (i.e., if its IP address does not change over time, or if it changes slowly enough so that its movements can be tracked in real-time), a DNSBL operator could take active countermeasures.

Distributed Reconnaissance

Referring back to FIG. 10, in 1015, distributed reconnaissance is performed. In distributed reconnaissance, each bot performs reconnaissance on behalf of other bots either in the same botnet or in other botnets. This is done because single host third-party reconnaissance can be easily subject to detection. To remain more stealthy, and to distribute the workload of performing DNSBL reconnaissance, botmasters may distribute lookups across the botnet itself, having bots perform distributed reconnaissance. In this case, the number of lookups by each bot is small and close to the number of lookups on the bot itself. Thus, the α value of a bot could be close to that of legitimate servers. Thus, an additional method can be used to detect bots performing distributed reconnaissance.

The temporal arrival pattern of queries at the DNSBL by hosts performing reconnaissance may differ from temporal characteristics of queries performed by legitimate hosts. With legitimate mail server's DNSBL look-ups, the look-ups are typically driven automatically when email arrives at the mail server and will thus arrive at a rate that mirrors the arrival rates of email. Distributed reconnaissance-based look-ups, on the other hand, will not reflect any realistic arrival patterns of legitimate email. In other words, the arrival rate of look-ups from a bot is not likely to be similar to the arrival rate of look-ups from a legitimate email server.

FIG. 13 illustrates the process of determining whether the arrival rate of look-ups from a source IP are similar to the arrival rate of look-ups from legitimate email servers, according to one embodiment of the invention. In 1305, a list of known or probable legitimate email servers that are using the DNSBL service is identified. This can be done, for example, as set forth below:

If the DNSBL is subscription-based or has access control, use a list of approved users (the email servers) to record the IP addresses that the servers use for accessing the DNSBL service. Enter these addresses into a list of Known Mail Server IPs.

If the DNSBL service allows anonymous access, monitor the source IPs of incoming look-up requests, and record a list of unique IP addresses (hereinafter "Probable Known Mail Server IPs"). For each IP address in the Probably Known Mail Server IPs list:

Connect to the IP address to see if the IP address is running on a known mail server. If a banner string is in the return message from the IP address, and its responses to a small set of SMTP commands, e.g. VRFY, HELO, EHLO, etc., match known types and formats of responses associated with a typical known mail server, then the IP address is very likely to be a legitimate email server, and in such a case, enter it into the list of Known Mail Server IPs.

Those of skill in the art will understand that other methods may be used to compile a list of known legitimate email servers. In 1310, for each of the known or probable legitimate email servers, its look-ups to DNSBL are observed, and its average look-up arrival rate λ_i for a time interval (say, a 10-minute interval) is derived. This can be done, for example, by using the following simple estimation method. For n intervals (say n is 6), for each interval, the number of look-ups

14

from the mail server, d_k are recorded. The average arrival rate of look-ups from the mail servers over n time intervals is simply:

$$\lambda_i = \frac{\sum_{k=1}^n d_k}{n}$$

where:

λ_i = the average look-up rate for time interval i

d_k = the number of lookups from the known mail server

k = the known mail server

n = the number of time intervals

In 1315, once the look-up arrival rates from the known mail servers are learned, the average look-up arrival rate λ' from a source IP (that is not a known legitimate email server or a known bot) can be analyzed over n time intervals

In 1320, if λ' is very different from each λ_i , i.e., $\lambda' - \lambda_i > t$ for all i 's, where t is a threshold, the source IP is considered a bot. The above process of measuring the arrival rates of the legitimate servers is repeated for every n time intervals. The comparison of the arrival rate from a source IP, λ' , with the normal values, λ_i 's, is performed using the λ' and λ_i 's computed over the same period in time.

FIG. 15 illustrates a method for constructing a DNSBL query graph, according to one embodiment of the invention. Referring to FIG. 15, in 1505 a set of DNSBL query logs is input. In 1510, the DNSBL queries are parsed to include only querier or queried IP addresses. In 1515, the DNSBL queries are then pruned to include only IP addresses which are present in a set B, which is a set of known bot IP addresses. In 1520, a graph G is a DNSBL query graph constructed using the input from 1505-1515. G illustrates all IP addresses that are querier or queried by the DNSBL pruned queries. Thus, G illustrates all suspect IP addresses that either queried, or were queried by the suspect IP addresses in set B. In 1525, to address the situation where both the querier or queried nodes from the DNSBL query set are members of B, a query graph extrapolation is performed. Here a second pass is made and edges are added if at least one of the endpoints of the edge (i.e., either querier or queried) is already present on the graph G.

FIG. 16 is an algorithm setting forth the method explained in FIG. 15, according to one embodiment of the invention. FIG. 12 sets forth a table of nodes, found utilizing the algorithm in FIG. 16, which has the highest out-degrees, and the number of hosts that are known spammers (appearing in a spam sinkhole).

In addition to finding bots that perform queries for other IP addresses, the above methods also lead to the identification of additional bots. This is because when a bot has been identified as performing queries for other IP addresses, the other machines being queried by the bot also have a reasonable likelihood of being bots.

The above methods could be used by a DNSBL operator to take countermeasures (sometimes called reconnaissance poisoning) towards reducing spam by providing inaccurate information for the reconnaissance queries. Examples of countermeasures include a DNSBL communicating to a botmaster that a bot was not listed in the DNSBL when in fact it was, causing the botmaster to send spam from IP addresses that victims would be able to more easily identify and block. As another example, a DNSBL could tell a botmaster that a bot was listed in the blacklist when in fact it was not, potentially causing the botmaster to abandon (or change the use of)

a machine that would likely be capable of successfully sending spam. The DNSBL could also be integrated with a system that performs bot detection heuristics, as shown in FIG. 14. FIG. 14 illustrates spamming bots and a C&C performing reconnaissance, attempting to get DNSBL information. Legitimate DNSBL lookups from a victim's computer are also being requested. A DNSBL responds to the bots, the C&C, and the legitimate computer, but the DNSBL may respond in different ways. For example, the DNSBL may tell the bot computers wrong information in response to their DNSBL requests in order to confuse the botnet, while returning correct information to legitimate servers.

In addition, a known reconnaissance query could be used to boost confidence that the IP address being queried is in fact also a spamming bot. Furthermore, DNSBL lookup traces would be combined with other passively collected network data, such as SMTP connection logs. For example, a DNSBL query executed from a mail server for some IP address that did not recently receive an SMTP connection attempt from that IP address also suggests reconnaissance activity.

DNS Cache Snooping

FIGS. 17-18 illustrate a technique to estimate the population of bots within a network through DNS cache inspection or snooping, according to one embodiment of the invention. DNS non-recursive queries (or resolution requests for domains that the DNS server is not authoritative for) are used to check the cache in a large number of DNS servers on the Internet to infer how many bots are present in the network served by each DNS server. DNS non-recursive queries instruct the DNS cache not to use recursion in finding a response to the query. Non-recursive queries indicate in the query that the party being queried should not contact any other parties if the queried party cannot answer the query. Recursive queries indicate that the party being queried can contact other parties if needed to answer the query.

In general, most domain names that are very popular, and thus used extensively, are older, well-known domains, such as google.com. Because of the nature of botnets, however, although they are new, they are also used extensively because bots in the botnet will query the botnet C&C machine name more frequently at the local Domain Name Server (LDNS), and hence, the resource record of the C&C machine name will appear more frequently in the DNS cache. Since non-recursive DNS queries used for DNS cache inspection do not alter the DNS cache (i.e., they do not interfere with the analysis of bot queries to the DNS), they can be used to infer the bot population in a given domain. Thus, when the majority of local DNS servers in the Internet are probed, a good estimate of the bot population in a botnet is found.

DNS cache inspection utilizes a TTL (time-to-live) value (illustrated in FIG. 18) of the resource record of a botnet C&C domain to get an accurate view of how long the resource record stays in the DNS cache. (Note that IP addresses change and/or the DNS server can only remember cache information for a certain amount of time.) When the resource record is saved in the cache, (e.g., as a result of the first DNS look up of the C&C domain from the network), it has a default TTL value, set by the authoritative DNS server. As time goes on, the TTL value decreases accordingly until the resource record is removed from the cache when the TTL value drops to zero. Referring to FIG. 18, three caching episodes are illustrated, each with a beginning point in time b1, b2, and b3, and an end point in time e1, e2, e3. The distance between caching episodes is described as T1, T2, etc. Thus, if we see many caching episodes (or "shark fins") on FIG. 18, we can determine that a large number of hosts are attempting to contact the C&C domain. If the C&C domain is a relatively new and

unknown domain, we can then surmise that the domain is used for malicious purposes (e.g., botnet coordination).

Referring to FIG. 17, one embodiment of a DNS cache inspection technique is as follows: In 1705, probes are done for open recursion, and open recursive servers are identified. Open recursive servers are servers that will perform recursive DNS lookups on behalf of queries originating outside of their network. In 1710, priority ranking of domains is performed. (This process is described in more detail later.) The output of 1705 and 1710 (which can be independent phases) is then used in a non-recursive query in 1715. In 1720, analysis is performed, including: (a) determining the relative ranking of botnet sizes, (b) estimating the number of infected individuals/bots within a botnet, and (c) assessing whether and to what extent a given network has infected computers. Since infections are dynamic, ongoing probes are needed. Thus, the analysis from 1720 can also be used to redo 1715 and prioritize the work performed in 1715.

Identifying Open Recursive Servers

Open recursive servers can be identified to, for example: (a) estimate botnet populations, (b) compare the relative sizes of botnets, and (c) determine if networks have botnet infections based on the inspection of open recursive DNS caches.

Open recursive DNS servers are DNS servers that respond to any user's recursive queries. Thus, even individuals outside of the network are permitted to use the open recursive DNS server. The cache of any DNS server stores mappings between domain names and IP addresses for a limited period of time, the TTL period, which is described in more detail above. The presence of a domain name in a DNS server's cache indicates that, within the last TTL period, a user had requested that domain. In most cases, the user using the DNS server is local to the network.

In 1705 of FIG. 17, networks are scanned for all DNS servers, and the networks identify the servers that are open recursive DNS servers. A DNS server (and thus, an open recursive DNS server) can be operated at almost any address within the IPv4 space (i.e., that portion not reserved for special use). We refer to this usable IPv4 address space as a "routable address".

To speed up the search for all DNS servers on the Internet, 1705 breaks up the routable space into organizational units. The intuition is that not all IPv4 addresses have the same probability of running a DNS server. Often, organizations run just a handful of DNS servers, or even just one. The discovery of a DNS server within an organizational unit diminishes (to a non-zero value) the chance that other addresses within the same organization's unit are also DNS servers.

1705 is explained in more detail in FIG. 19, according to one embodiment of the invention. In 1905, the IPv4 routable addresses (using, for example, Request for Comments (RFC) 3330) (note that an RFC is a document in which standards relating to the operation of the Internet are published) is organized into organizational units (using for example, RFC 1446). In 1910, for each organizational unit in 1905, the following calculations are performed to obtain the classless interdomain routing (CIDR) Priority Ranking Score ("CPRS"):

- a. For each DNS server known to exist in the organizational unit, add 1.0.
- b. For each IP address unit that has previously been seen to not run a DNS server, add 0.01.
- c. For each IP address unit for which no information is available, add 0.1.

In 1915, the organizational units are sorted in descending order according to their CPRS values.

Domain Ranking

1710 of the DNS cache inspection process (which can be independent of **1705**) produces a set of candidate domains. In other words, this phase generates a list of “suspect” domains that are likely botnet C&C domains. There are multiple technologies for deriving such a suspect list. For example, one can use DDNS or IRC monitoring to identify a list of C&C domains. Those of ordinary skill in the art will see that DDNS monitoring technologies can yield a list of botnet domains.

Cache Inspection

1715 of the DNS cache inspection process combines the outputs of **1705** and **1710**. For each domain identified in **1710**, a non-recursive query is made to each non-recursive DNS server identified in **1705**. Thus, for the top N entries (i.e., the N units with the lowest scores in **1915**), the following steps are performed to determine if the DNS server is open recursive:

a. A non-recursive query is sent to the DNS server for a newly registered domain name. This step is repeated with appropriate delays until the server returns an NXDOMAIN answer, meaning that no such domain exists.

b. A recursive query is then immediately sent to the DNS server for the same domain name used in the previous non-recursive query. If the answer returned by the DNS server is the correct resource record for the domain (instead of NXDOMAIN), the DNS server is designated as open recursive.

Determine Number of DNS Servers

Once an open recursive server is discovered, its cache can be queried to find the server’s IP address. Often the server’s IP address can be hard to discover because of server load balancing. Load balancing is when DNS servers are clustered into a farm, with a single external IP address. Requests are handed off (often in round-robin style) to an array of recursive DNS machines behind a single server or firewall. This is illustrated in FIG. 20. Each DNS machine maintains its own unique cache, but the DNS farm itself presents a single IP address to outside users. Thus, an inspection of the DNS cache state could come (randomly) from any of the machines behind the single load balancing server or firewall.

This problem is addressed by deducing the number of DNS machines in a DNS farm. Intuitively, multiple non-recursive inspection queries are issued, which discover differences in TTL periods for a given domain. This indirectly indicates the presence of a separate DNS cache, and the presence of more than one DNS server behind a given IP address.

FIG. 21 illustrates a procedure used to deduce the number of DNS servers behind a load balancing server or firewall, according to one embodiment of the present invention. For each open recursive DNS server (ORN), it is determined if the DNS service is behind a load balancing server or firewall and if so the number of servers is estimated as follows: In **2105**, the number of Assumed DNS Servers (or “ADS”) is set to “1”. In **2110**, an existing domain is recursively queried for, and the TTL response time is observed. This can be called the TTL response TTL_0 , and can be placed into a table of Known TTL Values (“KTV”). In **2115**, a period of w_1 , w_2 , and w_3 seconds is waited, where all values of w are less than all KTV entries. In **2120**, after w_1 , w_2 , w_3 seconds, another query is sent to the server. The corresponding TTL response times are observed and called TTL_1 , TTL_2 , and TTL_3 . In **2125**, if $w_1 + TTL_1$ does not equal any value already in KTV, then TTL_1 is entered into the KTV table, and the number of ADS’s is incremented by one. This is repeated for $w_2 + TTL_2$, and $w_3 + TTL_3$. In **2130**, it is determined if the ADS count has not been incremented. If not, in **2140**, the system is exited. If yes, steps **2120-2130** are repeated until the number of ADS’s does not increase.

Some load balancing is performed by a load balancing switch (often in hardware) that uses a hash of the 4-tuple of

the source destination ports and IP addresses to determine which DNS server to query. That is, queries will always reach the same DNS server if the queries originate from the same source IP and port. To accommodate this type of load balancing, a variation of the above steps can be performed. **2115** through **2135** can be performed on different machines with distinct source IPs. (This may also be executed on a single multihomed machine that has multiple IP addresses associated with the same machine and that can effectively act as multiple machines.) Thus, instead of starting three threads from a single source IP address, three machines may each start a single thread and each be responsible for querying the DNS server from a distinct source IP. One of the machines is elected to keep track of the ADS count. The distributed machines each wait for a separate wait period, w_1 , w_2 , and w_3 , per step **2115**. The distributed machines coordinate by reporting the outcome of the results in steps **2120-2130** to the machine keeping track of the ADS count.

If all DNS queries use only (stateless) UDP packets, the queries may all originate from the same machine, but forge the return address of three distinct machines programmed to listen for the traffic and forward the data to the machine keeping track of the ADS count.

Once the ADS count has been determined for a given DNS server, cache inspection can be performed according to the procedure in FIG. 22. In **2205**, each domain identified in **1710** is called a $Domain_S$. For each $Domain_S$, the DNS start of authority (SOA) is consulted for the TTL. This value is called TTL_{SOA} . In **2210**, for an ORN, x threads are created, where $x = ADS * 2$ (2 times the number of Assumed DNS Servers). The threads are synchronized to perform DNS queries simultaneously according to the following procedure. For $Domain_S$,

A master thread waits for half the TTL_{SOA} period, and then instructs the child threads to send their DNS queries. (Since there are twice as many queries as ADS, there is a high probability that each of the DNS servers will receive once of the queries.)

If any of the threads querying an ORN (an open recursive DNS server) reports the ORN not having a cache entry for $Domain_S$, repeat step (a) immediately.

If all of the threads reports that the ORN has a cache entry for $Domain_S$, the smallest returned TTL for all of the threads is called TTL_{min} , and all of the threads for $TTL_{min} - 1$ seconds sleep before waking to repeat step (a).

In **2215**, the above cycle, from **2210(a)** to **2210(c)**, builds a time series data set of $Domain_S$ with respect to an open recursive DNS server. This cycle repeats until $Domain_S$ is no longer of interest. This occurs when any of the following takes place:

a. $Domain_S$ is removed from the list of domains generated by **1710**. That is, $Domain_S$ is no longer of interest.

b. For a period of $x \cdot TTL_{SOA}$ consecutive periods, fewer than y recursive DNS servers identified in **1705** have any cache entries for $Domain_S$. That is, the botnet is old, no longer propagating, and has no significant infected population. In practice, the sum of the $x \cdot TTL_{SOA}$ period may total several weeks.

In **2220**, the cycle from steps **2210(a)** to **2210(c)** can also stop when the open recursive DNS server is no longer listed as open recursive by **1705** (i.e., the DNS server can no longer be queried).

Analysis

The analysis phase **1720** takes the cache observations from **1715**, and for each domain, performs population estimates. In one embodiment, the estimates are lower and upper bound calculations of the number of infected computers in a botnet.

For example, a botnet could be estimated to have between 10,000 and 15,000 infected computers. One assumption made is that the requests from all the bots in a network follow the same Poisson distribution with the same Poisson arrival rate. In a Poisson process, the time interval between two consecutive queries is exponentially distributed. We denote the exponential distribution rate as λ . Each cache gap time interval, T_i , ends with a new DNS query from one bot in the local network, and begins some time after the previous DNS query. Thus, in FIG. 18, the cache interval for the first bot's request occurs between b_1 and e_1 . The time interval T_1 measures the distance between the end of the first caching episode e_1 , and the start of the second b_2 .

As illustrated in FIG. 18, for a given domain, each name resolution (DNS query) by a bot triggers a caching event with a fresh TTL value that decays linearly over time. The time between any two caching episodes is designated T_i . The "memoryless" property of exponential distribution indicates that the cache gap time interval T_i follows the same exponential distribution with the same rate λ , no matter when the cache gap time interval begins. A function is said to be memoryless when the outcome of any input does not depend on prior inputs. All exponentially distributed random variables are memoryless. In the context of the DNS cache inspection, this means that the length of the current cache interval T_i does not depend on the length of the previous cache interval T_{i-1} .

Lower Bound Calculation.

A lower bound can be calculated on the estimated bot population. For the scenario depicted in the figure above, there was at least one query that triggered the cache episode from b_1 to e_1 . While there may have been more queries in each caching episode, each caching event from b_i to e_i represents at least a single query.

If λ_l is a lower bound (l) for the arrival rate, and T_i is the delta between two caching episodes, and M is the number of observations, for $M+1$ cache inspections, λ_l can be estimated as:

$$\frac{1}{\hat{\lambda}_l} = \sum_{i=1}^M \frac{T_i + TTL}{M} = TTL + \sum_{i=1}^M \frac{T_i}{M}$$

Using analysis of a bot (e.g., by tools for bot binary analysis), the DNS query rate λ can be obtained for each individual bot. Then from the above formula, the estimate of the bot population \hat{N}_l , in the network can be derived as follows:

$$\hat{N}_l = \frac{\hat{\lambda}_l}{\lambda}$$

Upper Bound Calculation.

During a caching period, there are no externally observable effects of bot DNS queries. In a pathological case, numerous queries could arrive just before the end of a caching episode, e_i . An upper bound can be calculated on the estimated bot population. Define λ_u as the upper bound estimate of the Poisson arrival rate. For the upper bound estimate, there are queries arriving between the times b_i and e_i . The time intervals T_i , however, represent periods of no arrivals, and can be treated as the sampled Poisson arrival time intervals of the underlying Poisson arrival process. It is fundamental that random, independent sample drawn from a Poisson process is itself a Poisson process, with the same arrival rate. This sampling is called the "Constructed Poisson" process.

For M observations, the estimated upper bound (u) arrival rate λ_u is:

$$\frac{1}{\hat{\lambda}_u} = \sum_{i=1}^M \frac{T_i}{M}$$

The population of victims needed to generate the upper bound arrival rate λ_u can therefore be estimated as:

$$\hat{N}_u = \frac{\hat{\lambda}_u}{\lambda}$$

CONCLUSION

While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example, and not limitation. It will be apparent to persons skilled in the relevant art(s) that various changes in form and detail can be made therein without departing from the spirit and scope of the present invention. In fact, after reading the above description, it will be apparent to one skilled in the relevant art(s) how to implement the invention in alternative embodiments. Thus, the present invention should not be limited by any of the above-described exemplary embodiments.

In addition, it should be understood that the figures and algorithms, which highlight the functionality and advantages of the present invention, are presented for example purposes only. The architecture of the present invention is sufficiently flexible and configurable, such that it may be utilized in ways other than that shown in the accompanying figures and algorithms.

Further, the purpose of the Abstract of the Disclosure is to enable the U.S. Patent and Trademark Office and the public generally, and especially the scientists, engineers and practitioners in the art who are not familiar with patent or legal terms or phraseology, to determine quickly from a cursory inspection the nature and essence of the technical disclosure of the application. The Abstract of the Disclosure is not intended to be limiting as to the scope of the present invention in any way.

What is claimed is:

1. A method of detecting a collection of compromised networks and/or computers, comprising:
 - performing processing associated with collecting Domain Name System (DNS) data, utilizing a detection system in communication with a database, the DNS data generated by a DNS server and/or similar device, wherein the DNS data comprises DNS queries, wherein the collected DNS data comprises DNS query rate information, and wherein the collecting DNS data from the DNS server comprises:
 - performing processing associated with identifying a command and control (C&C) computer in a first network, comprising:
 - performing processing associated with determining whether a computer has a suspicious DNS request rate, comprising: performing processing associated with calculating a canonical sub-level domain (SLD) request rate for a given SLD, wherein the canonical SLD request rate is calculated as the total number of requests to third level domains (3LDs)

21

present in the given SLD plus any request to the given SLD, and performing processing associated with determining whether the canonical SLD request rate of the given SLD significantly deviates from the mean of canonical request rates of SLDs; 5

when the DNS request rate is suspicious, performing processing associated with determining whether the DNS data has an exponential request rate comprising: performing processing associated with sorting DNS request rates per epoch, and performing 10 processing associated with determining whether there is exponential activity over a longer time epoch; and

when the DNS data has an exponential request rate, performing processing associated with identifying 15 the computer as the C&C computer; and

performing processing associated with recording an IP address and/or traffic information from a compromised computer when the compromised computer contacts 20 another computer;

performing processing associated with examining the collected DNS data relative to DNS data from known comprised and/or uncompromised computers; and

performing processing associated with determining an existence of the collection of compromised networks 25 and/or computers, and/or an identity of compromised networks and/or computers, based on the examination.

2. The method of claim 1, wherein collecting DNS data further comprises:

performing processing associated with replacing an IP 30 address of the C&C computer with an IP address of another computer, causing the compromised computer seeking to contact the C&C computer to be redirected to the other computer.

3. The method of claim 2, wherein the other computer is a 35 sinkhole computer.

4. The method of claim 1, further comprising:

performing processing associated with observing time zone and time of release information for the collected 40 data.

5. The method of claim 1, wherein determining the existence of the collection of compromised networks and/or computers is accomplished without contacting any networks or computers in the collection of compromised networks and/or 45 computers.

6. The method of claim 2, further comprising:

performing processing associated with isolating the collection of compromised networks and/or computers from its C&C computer, causing the collection of compromised 50 networks and/or computers to lose the ability to act as a coordinated group.

7. The method of claim 2, further comprising:

analyzing traffic from the compromised computer to the sinkhole computer to obtain information about a malware author.

8. The method of claim 4, further comprising:

utilizing time zone and time of release information to predict optimal release time information for an attack.

9. A system for detecting a collection of compromised networks and/or computers, comprising:

a computer, adapted to receive Domain Name System (DNS) data from a DNS server utilizing a detection system in communication with a database, the detection system configured for:

performing processing associated with collecting Domain 60 Name System (DNS) data, utilizing a detection system in communication with a database, the DNS data gener-

22

ated by a DNS server and/or similar device, wherein the DNS data comprises DNS queries, wherein the collected DNS data comprises DNS query rate information, and wherein the collecting DNS data from the DNS server comprises:

performing processing associated with identifying a command and control (C&C) computer in a first network, comprising:

performing processing associated with determining whether a computer has a suspicious DNS request rate, comprising: performing processing associated with calculating a canonical sub-level domain (SLD) request rate for a given SLD, wherein the canonical SLD request rate is calculated as the total number of requests to third level domains (3LDs) present in the given SLD plus any request to the given SLD, and performing processing associated with determining whether the canonical SLD request rate of the given SLD significantly deviates from the mean of canonical request rates of SLDs; 5

when the DNS request rate is suspicious, performing processing associated with determining whether the DNS data has an exponential request rate comprising: performing processing associated with sorting DNS request rates per epoch, and performing 10 processing associated with determining whether there is exponential activity over a longer time epoch; and

when the DNS data has an exponential request rate, performing processing associated with identifying 15 the computer as the C&C computer; and

performing processing associated with recording an IP address and/or traffic information from a compromised computer when the compromised computer contacts another computer;

performing processing associated with examining the collected DNS data relative to DNS data from known 20 comprised and/or uncompromised computers; and

performing processing associated with determining an existence of the collection of compromised networks and/or computers, and/or an identity of compromised networks and/or computers, based on the examination.

10. The system of claim 9, wherein collecting DNS data further comprises:

performing processing associated with replacing an IP 25 address of the C&C computer with an IP address of another computer, causing the compromised computer seeking to contact the C&C computer to be redirected to the other computer.

11. The system of claim 10, wherein the other computer is a sinkhole computer.

12. The system of claim 10, wherein the computer is further 30 capable of:

performing processing associated with observing time zone and time of release information for the collected data.

13. The system of claim 9, wherein determining the existence of the collection of compromised networks and/or computers is accomplished without contacting any networks or computers in the collection of compromised networks or 35 computers.

14. The system of claim 10, wherein the detection system is further configured for:

performing processing associated with isolating the collection of compromised networks and/or computers from 40

23

its C&C computer, causing the collection of compromised networks and/or computers to lose the ability to act as a coordinated group.

15. The system of claim 10, wherein the detection system is further configured for:

performing processing associated with analyzing traffic from the compromised computer to the sinkhole computer to obtain information about a malware author.

16. The system of claim 12, wherein the detection system is further configured for:

performing processing associated with utilizing time zone and time of release information to predict optimal release time information for an attack.

17. The method of claim 1, wherein collecting DNS data comprises:

performing processing associated with determining whether a source Internet Protocol (IP) address performing reconnaissance belongs to a compromised computer, the source IP address looking up at least one subject IP addresses; and

when the source IP is known to belong to a compromised computer, performing processing associated with designating the at least one subject IP addresses as a compromised computer.

18. The method of claim 17, wherein determining whether the source IP address belongs to a compromised computer comprises:

performing processing associated with determining whether the source IP address is a known compromised computer utilizing a DNS-based Blackhole List (DNSBL) and/or another list of compromised computers.

19. The method of claim 18, wherein determining whether the source IP address belongs to a compromised computer comprises:

performing processing associated with determining whether the source IP address is also the subject IP address.

20. The method of claim 18, wherein determining whether the source IP address belongs to a compromised computer comprises:

performing processing associated with determining a look-up ratio for the source IP address, the look-up ratio comprising the number of IP addresses the source IP address queries divided by the number of IP addresses that issue a look-up for the source IP address; and

when the look-up ratio for the source IP address is high, designating the source IP address as a compromised computer.

21. The method of claim 18, wherein determining whether the source IP address belongs to a compromised computer comprises:

performing processing associated with determining a look-up ratio for the source IP address, the look-up ratio comprising the number of IP addresses the source IP queries divided by the number of IP addresses that issue a look-up for the source IP address;

when the look-up ratio for the source IP address is low, performing processing associated with determining whether the look-up arrival rate mirrors the email arrival rate; and

the look-up arrival rate does not mirror the email arrival rate, performing processing associated with designating the source IP address as a compromised computer.

22. The method of claim 21, wherein determining whether the look-up arrival rate mirrors the email arrival rate further comprises:

24

performing processing associated with identifying a list of known and/or probably legitimate IP addresses using a DNSBL service;

for each of the known and/or probably legitimate IP addresses, performing processing associated with determining its average look-up arrival rate;

performing processing associated with determining an average look-up arrival rate from the source IP address;

performing processing associated with comparing the average look-up rates of the known and/or probably legitimate IP addresses to the arrival rate from the source IP address; and

when the average look-up rates of the known and/or probably legitimate IP addresses differ significantly from the arrival rate from the source IP address, performing processing associated with designating the source IP address as a compromised computer.

23. The method of claim 22, wherein identifying a list of known IPs comprises:

when the DNSBL service has controlled access, performing processing associated with recording IP addresses of approved users.

24. The method of claim 22, wherein identifying a list of probably legitimate IPs comprises:

when the DNSBL service allows anonymous access, performing processing associated with monitoring the source IP addresses of incoming look-up requests, and recording these source IP addresses;

performing processing associated with connecting to the IP address to determine whether the IP address is running on a known server; and

when the IP address is running on a known server, performing processing associated with designating the IP address as probably legitimate.

25. The system of claim 9, wherein collecting DNS data comprises:

performing processing associated with determining whether a source Internet Protocol (IP) address performing reconnaissance belongs to a compromised computer, the source IP address looking up at least one subject IP addresses; and

when the source IP is known to belong to a compromised computer, performing processing associated with designating the at least one subject IP addresses as a compromised computer.

26. The system of claim 25, wherein determining whether the source IP address belongs to a compromised computer comprises:

performing processing associated with determining whether the source IP address is a known compromised computer utilizing DNSBL and/or another list of compromised computers.

27. The system of claim 25, wherein determining whether the source IP address belongs to a compromised computer comprises:

performing processing associated with determining whether the source IP address is also the subject IP address.

28. The system of claim 25, wherein determining whether the source IP address belongs to a compromised computer comprises:

performing processing associated with determining a look-up ratio for the source IP address, the look-up ratio comprising the number of IP addresses the source IP address queries divided by the number of IP addresses that issue a look-up for the source IP address; and

25

when the look-up ratio for the source IP address is high, performing processing associated with designating the source IP address as a compromised computer.

29. The system of claim 25, wherein determining whether the source IP address belongs to a compromised computer 5 comprises:

performing processing associated with determining a look-up ratio for the source IP address, the look-up ratio comprising the number of IP addresses the source IP queries divided by the number of IP addresses that issue 10 a look-up for the source IP address;

when the look-up ratio for the source IP address is low, performing processing associated with determining whether the look-up arrival rate mirrors the email arrival rate; and 15

when the look-up arrival rate does not mirror the email arrival rate, performing processing associated with designating the source IP address as a compromised computer.

30. The system of claim 29, wherein determining whether the look-up arrival rate mirrors the email arrival rate further 20 comprises:

performing processing associated with identifying a list of known and/or probably legitimate IP addresses using a DNSBL service; 25

for each of the known and/or probably legitimate IP addresses, performing processing associated with determining its average look-up arrival rate;

performing processing associated with determining an average look-up arrival rate from the source IP address; 30

performing processing associated with comparing the average look-up rates of the known and/or probably legitimate IP addresses to the arrival rate from the source IP address; and

when the average look-up rates of the known and/or probably legitimate IP addresses differ significantly from the arrival rate from the source IP address, performing processing associated with designating the source IP address as a compromised computer. 35

31. The system of claim 30, wherein identifying a list of known IPs comprises: 40

when the DNSBL service has controlled access, performing processing associated with recording IP addresses of approved users.

32. The system of claim 30, wherein identifying a list of probably legitimate IPs comprises: 45

when the DNSBL service allows anonymous access, performing processing associated with monitoring the source IP addresses of incoming look-up requests, and performing processing associated with recording these source IP addresses; 50

performing processing associated with connecting to the IP address to determine whether the IP address is running on a known server; and

when the IP address is running on a known server, performing processing associated with designating the IP address as probably legitimate. 55

33. The method of claim 1, wherein collecting DNS data comprises:

performing processing associated with identifying open recursive DNS servers; and 60

performing processing associated with priority ranking domain names.

34. The method of claim 33, wherein the determining comprises: 65

performing processing associated with utilizing the open recursive DNS servers and the priority-ranked domain

26

names to determine whether the open recursive DNS servers are compromised computers.

35. The method of claim 34, further comprising:

performing processing associated with ranking sizes of networks of compromised computers;

performing processing associated with estimating a number of compromised computers in a network;

performing processing associated with assessing to what extent a given network has compromised computers;

performing processing associated with determining a lower bound calculation of a compromised computer population; or

performing processing associated with determining an upper bound calculation of a compromised computer population; or

any combination thereof.

36. The method of claim 33, wherein identifying open recursive DNS servers comprises:

performing processing associated with organizing IPv4 routable addresses into units;

performing processing associated with determining a classless interdomain routing (CIDR) priority ranking score (CPRS) value for each unit utilizing DNS server information; 25

performing processing associated with sorting the units a list in descending order utilizing the CPRS value; and

performing processing associated with determining whether a DNS server is an open recursive DNS server for DNS servers in the top of the list.

37. The method of claim 36, wherein determining the CPRS value comprises:

performing processing associated with giving a value of 1.0 for each DNS server known to exist in the unit;

performing processing associated with giving a value of 0.01 for each IP address known to run on a DNS server; and

performing processing associated with giving a value of 0.1 for each IP address with no DNS server information.

38. The method of claim 34, wherein utilizing the open recursive DNS servers and the priority-ranked domains in a recursive query comprises:

performing processing associated with sending at least one non-recursive query to the DNS server for a newly registered domain until the DNS server returns an NXDOMAIN answer;

performing processing associated with immediately sending a recursive query to the DNS server for the newly registered domain; and

performing processing associated with designating the DNS server as open recursive when the answer returned by the DNS server is not NXDOMAIN.

39. The method of claim 36, further comprising performing processing associated with determining the number of DNS servers behind a load balancing server.

40. The system of claim 9, wherein collecting DNS data comprises:

performing processing associated with identifying open recursive DNS servers; and

performing processing associated with priority ranking domain names.

41. The system of claim 40, wherein the determining comprises:

performing processing associated with utilizing the open recursive DNS servers and the priority-ranked domain names to determine when the open recursive DNS servers are compromised computers.

27

42. The system of claim 41, further comprising:
 performing processing associated with ranking sizes of
 networks of compromised computers;
 performing processing associated with estimating a num-
 ber of compromised computers in a network; 5
 performing processing associated with assessing to what
 extent a given network has compromised computers;
 performing processing associated with determining a
 lower bound calculation of a compromised computer
 population; or 10
 performing processing associated with determining an
 upper bound calculation of a compromised computer
 population; or
 any combination thereof.

43. The system of claim 40, wherein identifying open 15
 recursive DNS servers comprises:

performing processing associated with organizing IPv4
 routable addresses into units;
 performing processing associated with determining a
 CPRS value for each unit utilizing DNS server informa- 20
 tion;
 performing processing associated with sorting the units a
 list in descending order utilizing the CPRS value; and
 performing processing associated with determining
 whether a DNS server is an open recursive DNS server 25
 for DNS servers in the top of the list.

44. The system of claim 43, wherein determining the CPRS
 value comprises:

28

performing processing associated with giving a value of
 1.0 for each DNS server known to exist in the unit;
 performing processing associated with giving a value of
 0.01 for each IP address known to run on a DNS server;
 and

performing processing associated with giving a value of
 0.1 for each IP address with no DNS server information.

45. The system of claim 41, wherein utilizing the open
 recursive DNS servers and the priority-ranked domains in a
 recursive query comprises: 10

performing processing associated with sending at least one
 non-recursive query to the DNS server for a newly reg-
 istered domain until the DNS server returns an NXDO-
 MAIN answer;

performing processing associated with immediately send- 15
 ing a recursive query to the DNS server for the newly
 registered domain; and

when the answer returned by the DNS server is not NXDO-
 MAIN, performing processing associated with designat-
 ing the DNS server as open recursive. 20

46. The system of claim 45, further comprising performing
 processing associated with determining the number of DNS
 servers behind a load balancing server.

47. The method of claim 1, wherein the DNS data is non-
 recursive. 25

48. The system of claim 9, wherein the DNS data is non-
 recursive.

* * * * *